Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 210 (2023) 105669
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Herding through booms and busts ☆

Edouard Schaal [a,*], Mathieu Taschereau-Dumouchel [b]

[a] CREI, ICREA, UPF, BSE and CEPR, Spain
[b] Cornell University, United States of America

Received 10 May 2022; final version received 16 March 2023; accepted 28 April 2023
Available online 10 May 2023

Abstract

This paper explores whether rational herding can generate endogenous aggregate fluctuations. We embed a tractable model of rational herding into a business cycle framework. In the model, technological innovations arrive with unknown qualities, and agents have dispersed information about how productive the technology really is. Rational investors decide whether to invest based on their private information and the investment behavior of others. Herd-driven boom-bust cycles arise endogenously in this environment when the technology is unproductive but investors' initial information is overly optimistic. Their overoptimism leads to high investment rates, which investors mistakenly attribute to good fundamentals, leading to a self-reinforcing pattern of higher optimism and higher investment until the economy reaches a peak, followed by a crash when agents ultimately realize their mistake. We calibrate the model to the U.S. economy and show that it can broadly explain boom-and-bust cycles like the dot-com bubble of the 1990s.

JEL classification: E32; D80

☆ We thank Vladimir Asriyan, Andy Atkeson, Olivier Coibion, Pablo Fajgelbaum, Axelle Ferriere, Jordi Gali, Wouter den Haan, Nir Jaimovich, Francesco Lippi, Claudio Michelacci, Luigi Paciello, Ricardo Reis, Victor Rios-Rull, Nancy Stokey, Gilles Saint-Paul, Victoria Vanasco, Jaume Ventura, Jon Vogel, Pierre-Olivier Weill and Yi Wen for helpful comments. We thank participants at various seminars and conferences. Edouard Schaal acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563 and CEX2019-000915-S) and through the program Proyecto I+D Excelencia 2017 (ECO2017-87827-P-AEI/FEDER, UE), and support from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393). This paper was previously circulated under the title "Herding Cycles".

* Corresponding author.
E-mail addresses: eschaal@crei.cat (E. Schaal), mt763@cornell.edu (M. Taschereau-Dumouchel).

https://doi.org/10.1016/j.jet.2023.105669
0022-0531/

Keywords: Endogenous business cycles; Incomplete information; Herding

1. Introduction

Business cycle history is replete with examples in which new technologies led to periods of massive investment that ended in severe economic downturns. One salient example is the 1990s boom in information technologies that culminated in the stock market crash of 2001 ("dot-com bubble"). While the internet had been invented years earlier to connect academic and military networks, its commercial potential only became clear in the 1990s, when extreme enthusiasm for the new technology led to large investments in communication networks, software, and IT equipment. The high volume of investment and rising valuations of IT companies initially seemed to validate an optimistic outlook, but a crash eventually followed as some of the expected returns failed to materialize.[1] While the deep drivers that caused this sequence of events are still debated, a common view is that shifts in expectations played a key role in shaping the dot-com boom-bust cycle.

The idea that expectations contribute to aggregate economic fluctuations has a long tradition in macroeconomics. In seminal work, Pigou (1927) emphasized the importance of beliefs in shaping the business cycle. In his view, booms can be caused by waves of optimism among business executives, and crashes arise when their lofty expectations turn out to be mistaken. This hypothesis has been extensively studied in modern business cycle theory by the news-driven business cycle literature, pioneered by Beaudry and Portier (2004).[2] According to this view, agents receive news about future productivity, which sometimes turn out to be false. Boom-bust cycles arise after an initial sequence of positive news is later contradicted by experience.

These theories, however, remain mostly silent on the technological, social and psychological factors that drive the evolution of beliefs. In most of these studies, the belief process obeys an exogenous law of motion, and boom-bust cycles occur after a specific sequence of shocks—first positive, then negative. In other words, a large part of these cycles remains attributed to unexplained factors, precluding a deeper understanding of the key determinants of business cycles. What explains that beliefs follow this particular—and perhaps systematic—pattern which evolves from a phase of rising optimism to all-out pessimism? Is the growing optimism during the boom the consequence of luck, or the result of particular interactions between investors that lead to instability and inefficiencies? What causes precipitate the economy into a bust? Answering these questions is essential for our understanding of business cycles and for the design of stabilization policies.

This paper proposes one potential unifying explanation by exploring how herding in investors' behavior can generate a full macroeconomic boom-bust cycle without relying on an exogenous sequence of shocks. In our theory, entrepreneurs infer the quality of their investment opportunities by observing the decisions of others and can be tempted to invest when they see their

[1] Other boom-bust episodes follow similar patterns. For instance, the Roaring Twenties, a period of massive economic growth fueled by technological innovations in many sectors such as car manufacturing, communication, aviation and the chemical industry, ended in the Great Depression. Xiong (2013) documents several instances of boom-bust episodes that follow the introduction of new technologies. Arif and Lee (2014) document that aggregate investment tends to peak during periods of optimisms, and that these periods are followed by lower equity returns.

[2] See Beaudry and Portier (2014) for an overview of the empirical and theoretical research supporting the news view of the business cycle.

competitors expand their operations. The introduction of a new technology of uncertain quality can trigger a slow-rising boom followed by a sudden crash, in line with the experience of the dot-com era. In the boom phase, the initial optimism of investors translates into high levels of aggregate investment, and high investment, in turn, leads to further increases in optimism. This self-reinforcing process can fuel a long-lasting expansion of the economy, which comes to an end when new observations no longer support an optimistic view of the technology. Agents stop investing and the economy rapidly collapses. *Herding* thus offers a potential explanation for the emergence of technology-driven boom-bust cycles.

Our theory captures these ideas as follows. In the model, random technological innovations arrive over time and rational entrepreneurs decide whether to adopt the technology or not. The payoff from adoption is initially unknown, and agents use all available information to update their beliefs about the fundamental value of the technology. Information comes from both public and private sources. Importantly, to capture the idea that agents collect information from similar sources (news media, market reports, etc.), we assume that private signals feature some common noise. This assumption is key as it allows the distribution of beliefs across entrepreneurs to vary for reasons unrelated to the fundamental value of the technology. Entrepreneurs do not initially know the extent of that bias but progressively learn about it.

Agents also receive public signals. First, they learn by observing an exogenous public signal which stands in for the general information provided by public sources. They also learn from endogenous market outcomes such as aggregate quantities or prices. In the model, this amounts to observing, with some noise, the mass of agents who adopt the new technology. As the individual adoption decisions reflect the private information of the agents, this public signal operates as a *social learning* channel by aggregating, in a non-linear fashion, some of the information dispersed across agents.

How agents interpret this public signal is key for the emergence of boom-bust cycles. Such cycles are caused in our model by what we refer to as "false-positives": bad realizations of the technology fundamental that are accompanied by unusually large and positive realizations of the common noise. False-positives may thus capture situations in which, for instance, excessively promising benchmark tests are widely advertised upon the introduction of the technology and lead to overly optimistic beliefs.

When observing the large rate of adoption induced by such false-positive shocks, agents infer that private signals are positive. These signals, in turn, can be positive either because the fundamental value of the technology is good, or because the common noise component of the private signals is high. Entrepreneurs cannot tell these stories apart, but if false-positive shocks are relatively rare, the high adoption rate is initially attributed to a high-value technology, whose posterior likelihood rises. More optimistic beliefs lead to further aggregate adoption next period, which, in turn, leads to even more positive beliefs about the fundamental and so on. It is in that sense that our model displays a form of *herding*: agents mimic the behavior of others and sometimes mistakenly follow the herd into an adoption boom, meanwhile a shrinking measure of agents use their private information to go against the crowd. Through this positive feedback loop, the arrival of a low-value technology can create a long-lasting boom as entrepreneurs are fooled by the initial adoption craze.

But agents are rational and understand the possibility that they can sometimes be mistaken in their assessment of the true state of the world. As a result, they keep track of the probability of being in a false-positive state, which appears increasingly likely over time, as adoption keeps falling short of the most optimistic predictions. At some point, the most pessimistic agents stop adopting the technology and aggregate adoption no longer supports a high-productivity scenario.

This leads to a reversal in beliefs and a collapse in new technology adoption. We provide formal conditions under which these boom-and-bust episodes are guaranteed to arise in equilibrium.

A distinguishing feature of our approach is that the boom-and-bust cycle emerges *endogenously*. Standard practice in modern business cycle analysis often treats the booms and the busts as separate episodes, both driven by their own sequence of exogenous shocks. In contrast, our model generates an endogenous boom-and-bust cycle out of the single impulse shock that is the arrival of the new technology.[3] The crash, in particular, is not triggered by an exogenous shock but arises endogenously through the natural evolution of beliefs. As a consequence, the properties of the bust can be affected by what happened during the preceding boom, and government policies can have a large impact: policy interventions may affect the duration and magnitude of the boom as well as the timing and depth of the bust. They may also determine whether or not a cycle is to take place at all. This feature is absent from most standard models of the business cycle.

In the model, the mass of investing agents is a nonlinear aggregator of the dispersed information. As a result, the amount of information that agents receive is endogenous and varies with the cycle, which opens the door to a form of *information cascades*. When the public signals received up to a certain date are very positive, most agents invest regardless of their private signals so that their private information is not encoded into the mass of adopters. As a result, the model is able to generate sustained booms, when massive adoption restricts the flow of information, and rapid busts when slight downturns encourage enough entrepreneurs to use their private information, which suddenly reveals more information on the true state of the world.

Due to this variable flow of public information, the model features an information externality: agents do not internalize how their private adoption decisions affect the flow of public information. We characterize the solution of a social planning problem and show that the planner tends to lean against the wind by pushing for fewer agents to adopt the technology during booms and more during downturns so as to optimize the amount of information provided by aggregate adoption data.

To explore how the evolution of beliefs generated by our learning model can produce a general macroeconomic expansion followed by a recession, and to have a sense of the magnitude of the boom-bust cycles generated by the theory, we embed our main mechanism into a quantitative business cycle framework, which models the technology adoption decision of entrepreneurs after the arrival of a new technology. The model features two types of capital, "traditional" and "information technology" capital (IT), and we assume that the new technology is more intensive in IT capital. As in the basic model, social learning takes place as agents observe the measure of new-technology adopters.

We calibrate the model to match various moments of the data that relate to the dot-com period. In particular, we discipline the amount of private information—a key moment for our mechanism—using dispersion in forecasts from the Survey of Professional Forecasters (SPF). We also use data from the SPF to discipline investors' beliefs about the true value of the technology. Under our calibration, the model is able to generate a boom-bust cycle with positive

[3] By "endogenous", we mean that the entire boom-and-bust pattern is produced by the forces in the model. Our theory still relies on shocks, however, but only one-time shocks and does not rely on a particular sequence of positive then negative shocks. This approach is different from other theories of endogenous business cycles that generate deterministic periodic or chaotic dynamics (see Boldrin and Woodford (1990); Benhabib (1992); Guesnerie and Woodford (1992) for surveys).

comovements in consumption, investment, hours worked and output. The overinvestment into IT capital during the boom period and the negative wealth effect associated to its sudden loss of value cause the economy to contract significantly when beliefs collapse as agents realize that resources were misallocated. Overall, our results suggest that rational herding among economic agents can generate realistic fluctuations in macroeconomic aggregates.

Literature review

This paper builds on a long historical tradition in macroeconomics. The view that business cycles are shaped by expectations dates back at least to Pigou (1927), who also suggested a role of herding among investors.[4] While ignoring the financial aspects of booms and busts, our paper also echoes parts of the narrative that describes the behavioral and psychological causes of cycles in Minsky (1977) and Kindleberger (1978) after an initial "displacement" (e.g., the introduction of a new technology).

Our paper is closely related to the literature on news or noise-driven business cycles (Beaudry and Portier, 2004; Lorenzoni, 2009; Jaimovich and Rebelo, 2009). Indeed, our model shares the view that boom-bust cycles may be due to false-positives. In the news-shock literature, beliefs are driven by the exogenous release of news at fixed dates. In contrast, in our approach, the rise and fall in beliefs are *endogenously* driven by model forces, allowing us to explore the model's unique predictions on the frequency and timing of such cycles, and providing a greater role for stabilization policies. In addition, the news literature does not consider the role of herding in driving fluctuations, which is essential for our results: in our model, the gradually rising boom is the sole product of a positive feedback loop between investors under social learning, and the timing of the bust is determined by the time when the positive feedback disappears.

Christiano et al. (2008) consider the interaction of monetary policy and boom-bust cycles driven by news shocks. Closer to our work, Benhima (2019) builds a new-Keynesian model with dispersed private information in which news shocks released at fixed exogenous dates can create boom-bust episodes. In a recent paper, Angeletos et al. (2022) show that common noise shocks can generate excessive fluctuations in a model of start-up financing through complementary interactions between the entrepreneurial sector and financial markets. As in our setup, lean-against-the-wind policies can be beneficial. Burnside et al. (2016) propose an epidemiology-based model in which the transmission of optimistic beliefs in a population about the housing market can create a boom and bust. In a similar epidemiology-based variant of their framework, Angeletos and La'O (2013) propose a model in which a sentiment shock produces a hump-shaped cycle as it propagates through the population. Gorton and Ordoñez (2019) build a model in which the evolution of beliefs about the quality of collateral can lead to endogenous cycles. Goldstein et al. (2013) propose a model in which learning from prices creates strategic complementarities between financial markets and capital providers that give rise to trading frenzies.

Our paper also relates to the original work on herding and information cascades by Banerjee (1992), Bikhchandani et al. (1992) and Chamley (2004). It further relates to Avery and Zemsky (1998), who study herding in financial markets and introduce multidimensional uncertainty to

[4] In *Industrial Fluctuations* (1927), Pigou states that "the varying expectations of business men [...] and not anything else, constitute the immediate and direct causes or antecedents of industrial fluctuation". He emphasized the importance of the herding process: "the pioneers, who thus undertake and expand enterprises, at once fill a social need and lay treasure for themselves. Gradually, as no disaster happens to them, other less bold spirits follow their example; then others and yet other."

allow for information cascades. Our model differs from these traditional models of herding in several dimensions. First, in previous herding models, agents make decisions sequentially and the dynamics of the model are governed by the gradual observation of these individual decisions. Both features do not sit well with standard macroeconomic models. In our setup, instead, agents act simultaneously and learn by observing aggregates, which allows for a smoother integration of herding into macroeconomic frameworks. Second, the source of agents' confusion is different. In traditional herding models, people are confused between the fundamental return and the idiosyncratic shocks that stem from a particular ordering of the investors. As a consequence, boom-bust cycles arise only for specific sequences of idiosyncratic shocks. In our model instead, agents are confused between the fundamental and the common noise, which are drawn once and for all. Boom-bust patterns emerge endogenously through the natural evolution of beliefs and without any timing assumptions about shocks. This distinction with the existing literature is crucial to generate endogenous cycles.

To our knowledge, Loisel et al. (2012) is the only other macroeconomic model with herding. Their paper presents a simple general equilibrium model with overlapping generations of finitely-lived entrepreneurs who are endowed with private signals and must invest in a risky asset. As in traditional models of herding, entrepreneurs act sequentially and individual investment decisions are publicly observable. Our paper extends this approach by offering a novel herding model, based on contemporaneous decisions and the observation of aggregate actions, which, we believe, can be more easily integrated into traditional macroeconomic models.

Our learning model shares similarities with Vives (1997) who studies an environment in which agents with dispersed information learn by observing the average action across agents. Chapter 4 of Chamley (2004) briefly reviews a model in which privately informed agents learn from the average action. As in our model, the amount of information released by the public signal varies over the state space. As in Caplin and Leahy (1994) and Veldkamp (2005), the endogenous release of information in our model can generate sudden collapses in economic activity. Our work also contributes to a literature in which the aggregation of private information leads to nonlinear aggregate dynamics (Fajgelbaum et al., 2017). Straub and Ulbricht (2019) explore in a general setup the informativeness of nonlinear public signals and Straub and Ulbricht (2017) studies a particular application with financial constraints. None of these works consider the emergence of endogenous boom-bust cycles.

In contrast to our approach which maintains the assumption of rational expectations, a literature studies the emergence of boom-and-bust cycles in asset prices or in aggregate economic activity after departing from rationality or rational expectations. This includes the adaptive learning literature (Carceles-Poveda and Giannitsarou, 2008; Eusepi and Preston, 2011; Adam et al., 2017), the heterogeneous-belief literature with disagreement (Harrison and Kreps, 1978; Scheinkman and Xiong, 2003; Simsek, 2013) and, more recently, a literature that uses diagnostic expectations (Bordalo et al., 2021).

Our work also relates to a strand of literature that studies the role of bubbles in macroeconomic environments,[5] the literature on endogenous deterministic cycles,[6] a literature that views

[5] These include studies based on rational bubbles (Galí, 2014; Martin and Ventura, 2016; Asriyan et al., 2019; Guerrón-Quintana et al., 2020) as well as bubbles due to financial constraints and others (Kocherlakota, 1992; Miao and Wang, 2012; Barlevy, 2014; Hirano and Yanagawa, 2016).

[6] See, for instance, Grandmont (1985), Boldrin and Woodford (1990), Benhabib (1992), Benhabib et al. (2002) and Matsuyama (1999, 2013). More recent contributions include Beaudry et al. (2020), who provide empirical evidence in

endogenous cycles from the point of view of equilibrium indeterminacy and sunspots,[7] and a literature on learning from markets (Chen et al., 2007; Bond et al., 2012).

Section 2 introduces a simple learning model that conveys the intuition for the mechanism. The following section describes the forces at work in the model and discusses its welfare implications. Section 4 presents our business cycle model. We calibrate the model in Section 5 and show several empirical implications of the mechanism. The final section concludes.

2. Learning model

We first present our mechanism in a simplified dynamic game of technology adoption. This allows us to provide intuition for why social learning can lead to an endogenous herd-driven boom-bust cycle out of a single impulse shock. We also use this simplified model to derive analytical results and discuss the policy implications.

Notation

In what follows, whenever $F^x(\tilde{x}) = Pr(x \leq \tilde{x})$ denotes the cumulative distribution function (CDF) of some random variable x, f^x refers to its associated probability density function and $\overline{F}^x$ to its complementary CDF, $\overline{F}^x(\tilde{x}) = Pr(x > \tilde{x})$.

2.1. Environment

Time is discrete and goes on forever, $t = 0, 1, \ldots$. The economy is populated by a unit measure of entrepreneurs indexed by $j \in [0, 1]$. Entrepreneurs are risk-neutral and discount future consumption at rate $0 < \beta < 1$. Each entrepreneur faces a technology adoption choice and must decide whether to use an *old* technology or adopt a *new* one. The old technology is known and provides a constant deterministic return A^o. The new technology, on the other hand, is characterized by an unknown stochastic return A_t^n. We assume that the new technology is not immediately productive. It initially delivers the same return as the old technology until it matures, which happens with some fixed per-period probability $\lambda > 0$. After maturation, the true nature of the new technology is revealed and characterized by a constant return $\theta \in \{\theta_H, \theta_L\}$, $\theta_H > A^o > \theta_L$. We refer to θ as the technology *fundamental*. Maturation is a one-time event and all uncertainty is resolved afterwards. In other words,

$$A_t^n = \begin{cases} A^o & \text{before maturation} \\ \theta & \text{after maturation.} \end{cases}$$

Adopting the new technology is costless and, every period, agents must decide whether to use the new technology ($i_{jt} = 1$) or not ($i_{jt} = 0$). We assume that a fraction μ of agents are "noise entrepreneurs", that is, they are clueless regarding technological adoption and behave randomly.

favor of endogenous cycles. In contrast to our setup, models in this literature feature deterministic limit cycles that can be periodic or chaotic.

[7] This literature includes Benhabib and Farmer (1994) and Wen (1998). More recent contributions include Benhabib et al. (2015), Kaplan and Menzio (2016), Eeckhout and Lindenlaub (2019) and Golosov and Menzio (2020). These studies typically feature multiple equilibria and aggregate fluctuations are due to shifts in expectation triggered by sunspot shocks. Our model, instead, features a unique equilibrium and boom-bust cycles result from agents' gradual learning about the technology and the common noise.

Specifically, we assume that a fraction ε_t of noise entrepreneurs adopt the new technology, where ε_t is *iid* distributed according to a CDF F^{ε}. The remaining $1-\mu$ entrepreneurs are rational and choose the best of the two technologies, based on public and private information. There is no cost of switching between technologies, so entrepreneur j solves

$$i_{jt} = \underset{i_{jt}\in\{0,1\}}{\operatorname{argmax}}\, i_{jt} E\left[A_t^n \mid \mathcal{I}_{jt}\right] + \left(1-i_{jt}\right) A^o,$$

where $\mathcal{I}_{jt}$ is its information set at time t.

2.2. *Information*

The technology fundamental θ is randomly drawn once and for all at date 0. We denote by p_0 the ex-ante probability that $\theta=\theta_H$. Entrepreneurs do not observe θ directly but receive various private and public signals about its true value.

Private signals

First, we assume that each agent receives a private signal s_j at date 0, upon the arrival of the new technology. Importantly, we allow these private signals to feature not only idiosyncratic noise but also common noise. This common noise might come, for instance, from sources of information shared by agents (mass media, internet) that may report noisy signals about the initial success of the technology (e.g., benchmark tests). Common noise is key to our mechanism as it introduces the possibility that the average belief about θ varies for reasons that are orthogonal to the true value of the fundamental.

Common noise is captured by the random variable ξ, distributed according to the CDF F^{ξ}. Formally, we assume that the private signal s_j of agent j is drawn from the CDF $F^s_{\theta+\xi}(s) = Pr\left(s_j \le s\right)$, where $\left\{F^s_x\right\}_{x\in I}$ is a family of distributions that admit the probability density functions $\left\{f^s_x\right\}_{x\in I}$. To prevent the possibility of trivial learning, we make the assumption that F^s_x has full support over $\mathbb{R}$, i.e., $f^s_{\theta+\xi}>0$ everywhere. Finally, in order to guarantee monotonicity in learning, we assume that the family $\left\{F^s_x\right\}_{x\in I}$ satisfies the *monotone likelihood ratio property* (MLRP). That is, for $x_1<x_2\in I$ and $s_1<s_2$, we have

$$\frac{f^s_{x_2}(s_2)}{f^s_{x_1}(s_2)} \ge \frac{f^s_{x_2}(s_1)}{f^s_{x_1}(s_1)}. \tag{MLRP}$$

Intuitively, the MLRP condition guarantees that a high signal s is more likely to be coming from a high realization of $x=\theta+\xi$. In other words, an entrepreneur observing a high private signal s_j becomes more optimistic and puts a higher probability on the value of the technology θ and the common noise ξ being high.

Example. In most of our examples, we will use *additive private signals* so that

$$s_j = \theta + \xi + v_j, \text{ with } v_j \sim \text{iid CDF } F^v. \tag{1}$$

Public signals

Before the new technology matures, entrepreneurs collect public information in addition to their initial private signal. We first assume that each period all agents observe a public signal

$S_t = \theta + u_t$, centered around θ with *iid* noise distributed according to CDF F^u and standard deviation σ_u, which stands in for all the information collected over time from exogenous public sources. After maturation, we assume that the fundamental θ is observed and that all uncertainty is resolved.

Second, and more importantly for our mechanism, we introduce a form of *social learning* by allowing entrepreneurs to observe an endogenous signal which partially aggregates the private information of agents. This is to capture the type of information that agents learn by observing aggregate quantities or prices, which result from the aggregation of individual decisions.[8] Specifically, we assume that entrepreneurs observe the total measure of adopters of the new technology m_t, which, in the presence of noise entrepreneurs, can be written as[9]

$$m_t = \underbrace{\int_0^{1-\mu} i_{jt} dj}_{\text{rational entrepreneurs}} + \underbrace{\mu \varepsilon_t}_{\text{noise entrepreneurs}}. \tag{2}$$

The presence of the noise ε_t is required in our setting to prevent agents form learning too quickly (or even immediately in some cases, as we discuss later).

In equilibrium, the decision to adopt i_{jt} is a nonlinear function of the agent's individual beliefs. In turn, these beliefs are a function of public information up to time t, $\{S_{t-1}, m_{t-1}, \dots, S_0, m_0\}$, and of the private signal s_j. As a result, since public information is shared and can be filtered out, m_t partially aggregates the private information across the population of entrepreneurs and therefore contains useful information regarding the fundamental θ and the common noise ξ.

As we explain in more details below, it is the presence of this endogenous signal that will allow *herding* to occur in our environment. As m_t aggregates the dispersed information of private agents, a particularly high draw of m_t will be interpreted as being indicative of a strong fundamental θ and a high common noise ξ. As a result, agents will update their beliefs in favor of these states, which might encourage technology adoption, increasing m_t further, and so on.

2.3. Belief characterization

There are two aggregate shocks in this economy: the fundamental θ and the common noise ξ. The beliefs of an individual entrepreneur j are described by a joint probability distribution that we denote by

$$\Lambda_{jt}\left(\tilde{\theta}, \tilde{\xi}\right) = Pr\left(\theta = \tilde{\theta}, \xi \in \left[\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}\right] \mid \mathcal{I}_{jt}\right),$$

in which we explicitly allow for ξ to take a continuum of values and where $\mathcal{I}_{jt}$ is agent j's information set at date t. Since entrepreneurs receive different private signals, we should in principle keep track of the whole distribution of beliefs in the economy (i.e., a distribution over distributions). Fortunately, the information structure is simple enough that the model lends itself to a useful simplification. As in Chapter 3 of Chamley (2004), it is enough to keep track of only

[8] In our full business cycle model from Section 4, observing aggregate quantities or prices will provide a public signal of the same form.

[9] Alternatively, the noise ε_t in (2) can be interpreted as pure observational noise or measurement error in this simplified learning model.

one set of time-varying beliefs, the *public beliefs* $\Lambda_t\left(\tilde{\theta},\tilde{\xi}\right)=Pr\left(\theta=\tilde{\theta},\xi\in\left[\tilde{\xi},\tilde{\xi}+d\tilde{\xi}\right]\mid\mathcal{I}_t\right)$. These public beliefs correspond to the beliefs of an outside observer who only has access to public information $\mathcal{I}_t$ at time t, which is the collection of past public signals on the technology and measures of entrepreneurs: $\mathcal{I}_t=\{S_{t-1},m_{t-1},\ldots,S_0,m_0\}$. In comparison to this outside observer, an entrepreneur's information set also includes the private signal s_j, so that $\mathcal{I}_{jt}=\mathcal{I}_t\cup\left\{s_j\right\}$. Entrepreneurs' individual beliefs can easily be recovered from public beliefs using Bayes' rule and the private signal s_j, according to

$$\Lambda_{jt}\left(\tilde{\theta},\tilde{\xi}\right)=\frac{\Lambda_t\left(\tilde{\theta},\tilde{\xi}\right)f^s_{\tilde{\theta}+\tilde{\xi}}\left(s_j\right)}{\int\Lambda_t\left(\theta,\xi\right)f^s_{\theta+\xi}\left(s_j\right)d\left(\theta,\xi\right)}. \tag{3}$$

This simplification comes from the fact that only public information evolves over time. Indeed, since the private signal distribution $f^s_{\theta+\xi}$ is constant and known up to the realization of θ and ξ, it is easy to recover the entire distribution of private beliefs across agents for a given combination of (θ,ξ) at any point in time. As a result, the only object whose evolution over time we need to keep track of is the public belief distribution Λ_t.

2.4. Timing and adoption decision

The timing is as follows. At date 0, the fundamental θ, the common noise component ξ and the private signals s_j are drawn once and for all. At date $t\geq 0$,

1. Agents choose whether to adopt the new technology or not based on their individual beliefs Λ_{jt};
2. If it has not matured yet, the new technology matures with probability λ;
3. If the new technology matures, θ is learned. Otherwise, agents observe $\{S_t,m_t\}$ and update their beliefs. The economy moves to the next period.

Since the adoption decision is a simple static problem, it can be characterized in an easy way. Agent j adopts the technology in period t if and only if

$$E\left[A^n_t\mid\mathcal{I}_{jt}\right]=\lambda\left[p_{jt}\theta_H+\left(1-p_{jt}\right)\theta_L\right]+(1-\lambda)A^o\geq A^o, \tag{4}$$

where we define

$$p_{jt}=Pr\left(\theta=\theta_H\mid\mathcal{I}_{jt}\right)=\int\Lambda_{jt}\left(\theta_H,\xi\right)d\xi \tag{5}$$

as the probability that j puts on being in the good-technology state θ_H, the adoption decision (4) is characterized by a cutoff rule p^* in the space of beliefs. That is, an agent adopts the technology if and only if $p_{jt}\geq p^*$ where p^* is the belief of the marginal adopter such that

$$p^*\theta_H+\left(1-p^*\right)\theta_L=A^o. \tag{6}$$

The total measure of adopting agents, including noise entrepreneurs, can then be expressed as

$$m_t=m^e\left(\Lambda_t,\theta,\xi\right)+\mu\varepsilon_t, \tag{7}$$

$$\text{where } m^e\left(\Lambda_t,\theta,\xi\right)=(1-\mu)\int\mathbb{I}\left(p_j\left(\Lambda_t,s_j\right)\geq p^*\right)f^s_{\theta+\xi}\left(s_j\right)ds_j. \tag{8}$$

The variable m^e is the measure of adopting agents among rational entrepreneurs in a given state of the world (θ, ξ). Importantly for what follows, m^e is an object that any agent in the economy can compute. To see this, note that since they know the structure of the model and the public beliefs, all agents agree on the cutoff p^*. Second, thanks to the dichotomy between public beliefs and the fixed distribution of private signals $f^s_{\theta+\xi}$, agents understand the mapping from (θ, ξ) to p_j and can therefore compute the *distribution* of beliefs p_j that would arise in a given state of the world (θ, ξ). This property is essential to tractably solve the inference problem from the endogenous public signal, to which we now turn.

2.5. Evolution of beliefs

After characterizing the adoption decision, we can now describe how beliefs are updated over time. Before maturation, each end of period brings two new public signals for agents to process: S_t and m_t. The updating of information with S_t is straightforward as it is a simple exogenous signal. Applying Bayes' rule, we define the interim beliefs at the end of the period as

$$\Lambda_{t|S_t}\left(\tilde{\theta}, \tilde{\xi}\right) = \frac{\Lambda_t\left(\tilde{\theta}, \tilde{\xi}\right) f^u\left(S_t - \tilde{\theta}\right)}{\int \Lambda_t(\theta, \xi) f^u(S_t - \theta)\, d(\theta, \xi)}. \tag{9}$$

We now turn to incorporating the information contained in m_t. Solving the inference problem from an endogenous signal like m_t can be complicated in general because individual decisions need to be inverted to back out their information content about θ and ξ. Fortunately, and as highlighted at the end of the previous section, the inference problem is greatly simplified in our environment since the measure of rational adopters m^e in every state of the world is a simple function of the public beliefs Λ_t (known by everyone) and of the true realization of (θ, ξ). Entrepreneurs solely differ in their assessment of the probability of each state (θ, ξ), encoded in Λ_{jt}, but there is no *infinite regress* problem arising from the necessity to forecast the beliefs of agents after any history of shocks. Because of the equilibrium structure of the signal (7), Bayes' rule gives us the simple updating equation

$$\Lambda_{t+1}\left(\tilde{\theta}, \tilde{\xi}\right) = \frac{\Lambda_{t|S_t}\left(\tilde{\theta}, \tilde{\xi}\right) f^\varepsilon\left(\left(m_t - m^e\left(\Lambda_t, \tilde{\theta}, \tilde{\xi}\right)\right)/\mu\right)}{\int \Lambda_{t|S_t}(\theta, \xi) f^\varepsilon\left((m_t - m^e(\Lambda_t, \theta, \xi))/\mu\right) d(\theta, \xi)}. \tag{10}$$

2.6. Equilibrium

We are now ready to define an equilibrium in this economy.

Definition 1. An equilibrium consists of history-contingent public beliefs Λ_t, a distribution of private beliefs $\left\{\Lambda_{jt}\right\}_{j\in[0,1]}$ and a measure of technology adopters m_t for all t, such that: 1) the distribution of private beliefs is derived from the public beliefs through (3) and (5); 2) the measure of adopters is consistent with entrepreneurs decisions under their private beliefs as in (7); and 3) the public beliefs follow the laws of motion (9)–(10).

With that definition in hand, the following proposition characterizes the set of equilibria.

Proposition 1. *There exists a unique equilibrium.*

Proof. All proofs are in Appendix A.2. □

The proof of the proposition is straightforward. It shows that from a given distribution of public beliefs Λ_t, there is a unique mapping, given the realization of the shocks, to next period's public beliefs Λ_t. Starting from the initial Λ_0 we can therefore reconstruct the unique equilibrium sequence $\{\Lambda_0, \Lambda_1, \dots\}$. All other equilibrium quantities such as the measure of adopters and the distribution of private beliefs can then be reconstructed from the public beliefs in a unique way.

3. Endogenous booms and busts

We are now fully equipped to analyze the dynamics implied by the model. To simplify the exposition, we focus on a simple special case that conveys the intuition about the emergence of i) a smooth form of information cascades and ii) endogenous booms and busts. We then show that these results extend to a more general setup. We also discuss the welfare properties of the model.

3.1. The 3-state model

To simplify the exposition, we temporarily make the simplifying assumption that the pair (θ, ξ) can only take three different values, the minimal number of states required for endogenous boom-bust cycles to emerge in our model. Specifically, we assume[10]

$$(\theta, \xi) \in \left\{(\theta_L, 0), (\theta_H, 0), \left(\theta_L, \overline{\xi}\right)\right\} \text{ with } \theta_L < \theta_L + \overline{\xi} < \theta_H.$$

We refer to $(\theta_L, 0)$ as the *bad-technology* state, $(\theta_H, 0)$ as the *good-technology* state and $\left(\theta_L, \overline{\xi}\right)$ as the *false-positive* state. The latter is the state of interest as it is the one that will trigger a boom-and-bust cycle by having entrepreneurs mistakenly assess the technology to be of high quality before later realizing their mistake.

Having only three states reduces the number of state variables required to keep track of the belief distribution Λ_t. Public beliefs are now summarized by the two variables

$$p_t \equiv \Lambda_t (\theta_H, 0) \text{ and } q_t \equiv \Lambda_t \left(\theta_L, \overline{\xi}\right),$$

and the corresponding updating rules can be found in Appendix A.1.

We now establish a first result. Under our assumptions, the individual belief about the probability of the good technology, $p_{jt} = \Lambda_{jt} (\theta_H, 0)$, is increasing in the private signal s_j. As a result, the adoption decision can be further characterized by a cutoff rule $s^* (p_t, q_t)$ in terms of private signals, which simplifies the expression of the measure of rational adopters m^e as the following lemma shows.

Lemma 1. *In the three-state model, the optimal adoption strategy is characterized by a cutoff rule in the private signal $s^* (p_t, q_t)$, such that $p_j (p_t, q_t, s^* (p_t, q_t)) = p^*$, that is decreasing in p_t. That is, an agent adopts the technology if and only if $s_j \geq s^* (p_t, q_t)$. The measure of rational adopters is given by*

$$m^e (p_t, q_t, \theta, \xi) = (1 - \mu) \overline{F}^s_{\theta+\xi} \left(s^* (p_t, q_t)\right).$$

[10] We only consider the case $\theta_L < \theta_L + \overline{\xi} < \theta_H$ because it ensures that beliefs about the good state, p_{jt}, are non-decreasing in the signal s_j. This corresponds to the more general case from Section 3.2 where F^s satisfies the MLRP condition and ξ is normally distributed. In the case $\theta_L + \overline{\xi} > \theta_H$, beliefs are non-monotonic in s_j and other phenomena can be observed.

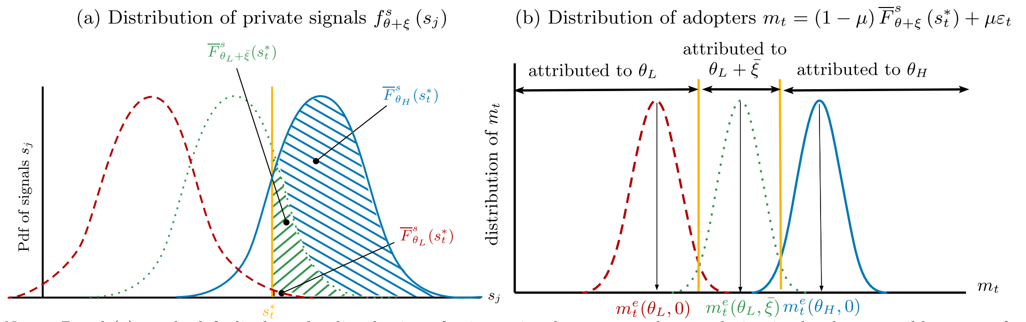

Notes: Panel (a) on the left displays the distribution of private signals s_j across the population in the three possible states of the world along with the corresponding expected measures of adopters $m^e_t = (1-\mu)\overline{F}^s_{\theta+\xi}(s^*(p_t, q_t))$, for some public beliefs (p_t, q_t). The dashed line represents the bad-technology state, the dotted line the false-positive state and the continuous line the good-technology state. Panel (b) on the right shows the distribution of $m_t = m^e_t + \mu\varepsilon_t$ in the three states of the world assuming some Gaussian-like distribution F^ε with mean 0 and variance σ^2_ε.

Fig. 1. Private beliefs and expected measure of adopters.

Learning from m_t

To develop intuition on the way agents learn from the measure of adopters, we propose an example in Fig. 1. Panel (a) displays the distribution of private signals s_j in the three states of the world. Due to the MLRP assumption, the three distributions are ordered in the sense of first-order stochastic dominance. The measure of rational adopters m^e is represented as the mass of agents located to the right of the cutoff s^*_t. We can see that m^e is small in the bad-technology state $(\theta_L, 0)$ (dashed line), that agents expect more adoption in the false-positive state $(\theta_L, \overline{\xi})$ (dotted line), and that it is at its largest in the good-technology state $(\theta_H, 0)$ (continuous line).

The three measures m^e being computed, we then present in panel (b) the three potential distributions of m_t in the three states of the world assuming that the noise ε is normally distributed with mean 0. As the graph illustrates, agents expect very different distributions of adopters m_t, each centered on their expected value m^e in the different states of the world (θ, ξ). We can split the m_t-space into three regions that indicate which state is attributed more probability after observing m_t. For instance, for low m_t the likelihood of the state θ_L is greater than that of the other states, so information updating will attribute it a higher probability. The two other states, $(\theta_L, \overline{\xi})$ and $(\theta_H, 0)$, have their own higher likelihood region that are also represented on the graph. Importantly for the emergence of boom-and-bust cycles, beliefs about the good state tend to increase after observing high realizations of m_t. It is in that sense that the model displays a form of "herding": agents become more optimistic (resp. pessimistic) after seeing high (resp. low) patterns of adoption, leading them to make inefficient adoption decisions, as we will see in our welfare analysis.

Signal-to-noise ratio and smooth information cascades

In the traditional herding literature (Banerjee, 1992; Bikhchandani et al., 1992), information cascades arise when public beliefs are so extreme (p_t extremely high or low, because of a particular history of public signals), that agents disregard their own private information. That is, agents adopt (or not) the technology no matter what their private information is. As a result, observing previous entrepreneurs' decisions becomes uninformative and the economy may end up being stuck in a situation with mistaken beliefs forever.

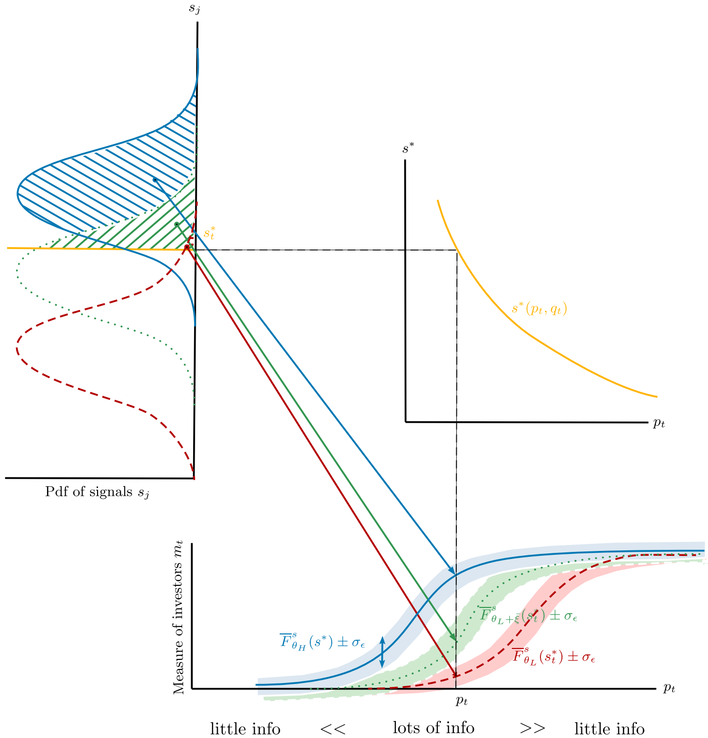

Notes: The top-left panel displays the distribution of private signals in the three states of the world along with the measure of rational technology adopters $m_t^e = (1-\mu)\,\overline{F}^s_{\theta+\xi}(s_t^*)$, as previously represented in Figure 1 rotated by 90°. The top-right panel displays an example of equilibrium threshold $s^*(p_t, q_t)$ as a function of public belief p_t. The bottom right panel shows how the measure of adopters $m_t = m_t^e + \mu\varepsilon$ varies with public belief p_t, keeping q_t constant in the background, under the assumption that $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon)$. The mean m_t^e is represented with a continuous line and the corresponding ±1-standard deviation $\mu\sigma_\varepsilon$ error bands with dashed lines.

Fig. 2. Measure of technology adopters m_t as a function of public belief p_t.

Because social learning takes place through the observation of the continuous variable m_t, rather than the sequence of binary decisions by previous entrepreneurs, the emergence of information cascades is somewhat different in our setup. We show nonetheless that a similar form of "smooth" information cascades may arise depending on assumptions about the distributions of signals.

The bottom-right panel of Fig. 2 represents how the measure of adopters m_t varies in expectation, along with its ±1-standard deviation error bands, as a function of the public belief p_t, holding q_t constant. These curves are drawn by first connecting a given level of p_t in the bottom-right panel to the equilibrium signal threshold $s^*(p_t, q_t)$ (upper-right panel), itself connected to the upper-left panel which shows how the measures $m^e = (1-\mu)\,\overline{F}^s_{\theta+\xi}(s^*)$ vary with the cutoff s^*. As the bottom panel shows, the expected measure of rational adopters m^e is a monotonic transformation of the CDF $F^s_{\theta+\xi}$ in the three different states.

The key feature to take away from this graph is that the *signal-to-noise* ratio in m_t varies nonmonotonically with the public beliefs. For intermediate values of p_t, the three expected measures m^e are far apart so that despite the noise ε_t, observing m_t is highly informative about the underlying state $\theta + \xi$ (i.e., the signal-to-noise ratio is high). For p_t large (resp. small), almost all (resp. no) agents adopt the technology, the three measures converge to $\lim_{s^* \to -\infty} \overline{F}^s_{\theta+\xi}(s^*) = 1$ (resp. 0 when $s^* \to \infty$), so that the signal m_t is dominated by noise and becomes uninformative about the underlying fundamentals (i.e., the signal-to-noise ratio is low). Note that this result is not an artifact of specific distributions or functional forms but is instead a general feature of the model as long as s^* varies sufficiently on the support of $\overline{F}^s_{\theta+\xi}$.[11]

The model offers a smooth analog to informational cascades when the equilibrium s^* reaches the extreme regions of the state space where learning is slow. Suppose for instance that public beliefs are optimistic (p_t high) so that s^* is very low. In such a situation, almost all agents act in the same way and adopt the new technology. Only few agents use their private information to "go against the crowd" and do not adopt: the most pessimistic ones that have received particularly low private signals. Unfortunately, their measure is so small that they are hard to detect when looking at the aggregate adoption patterns. As a result, markets are nearly uninformative and beliefs can remain wrong for an extended period of time. The main difference with traditional herding models is that, under the assumption that private signals have full unbounded support, the information flow is never exactly 0 so that there is always some learning taking place through m_t and S_t. Such a smooth form of information cascades is of interest to us for two reasons: i) it explains why the economy may remain for an extended period of time in the booming region, where agents understand that they could be wrong in their assessment of the true state of the world but adopt the technology nonetheless, ii) it opens the door to the economy endogenously exiting the information cascade and crashing when some threshold in beliefs is reached, as we will now describe.

Endogenous boom-and-bust cycle

We now turn to boom-and-bust cycles, which we define as a sequence of rising then declining public beliefs p_t about the good state. We present simulations of the model to illustrate its ability to generate endogenous boom-bust patterns out of a single impulse shock. We do not attempt to make a realistic calibration but merely pick parameters so as to highlight the model's properties. We will examine later under what general conditions one should expect the boom-bust cycles to occur.

We present the impulse responses of the measure of adopters (m_t) and the public beliefs (p_t, q_t), keeping all other shocks to their mean levels (e.g., ε_t, $u_t = 0$), when the economy is in the false-positive state $(\theta, \xi) = \left(\theta_L, \overline{\xi}\right)$, the case of interest for our purpose.[12]

Figs. 3 and 4 present two examples of endogenous boom-and-bust patterns that may arise in the model, depending on whether or not the economy falls into an information cascade. In both examples, the emergence of boom-and-bust patterns hinges on three key assumptions: (i) $\theta_L + \overline{\xi}$ needs to be sufficiently close to θ_H, so that the two states are hard to distinguish; (ii) S_t is not

[11] Fig. 2 may give the wrong impression that the nonmonotonicity result highly depends on the sigmoidal shape of the CDFs $F^s_{\theta+\xi}$. While the regions with higher signal-to-noise ratio may change with the distribution, a robust prediction for any distribution is that the measure m_t is less informative for extreme public beliefs, when agents herd on the same action, since $\overline{F}^s_{\theta+\xi}\left(s^*\right) \to 1$ (resp. $\overline{F}^s_{\theta+\xi}\left(s^*\right) \to 0$) when p_t gets close to 1 (resp. 0) and the cutoff s^* goes to minus infinity (resp. infinity) for any signal distribution.

[12] Figs. 9 and 10 in the appendix show the economy's response to the good-technology and bad-technology states. With our parametrization, these cases are relatively uninteresting: learning is fairly quick, and the dynamics are close to the full information case.

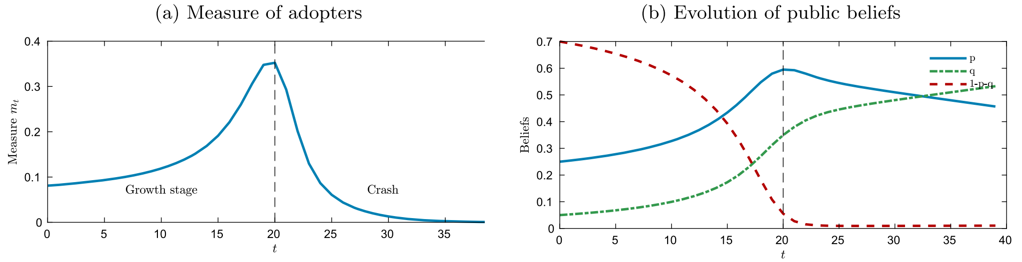

Notes: The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.53$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian: $F^s_{\theta+\xi} \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ and $F^u \sim \mathcal{N}(0, \sigma_u)$ with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 3$.

Fig. 3. Slow boom, sudden crash.

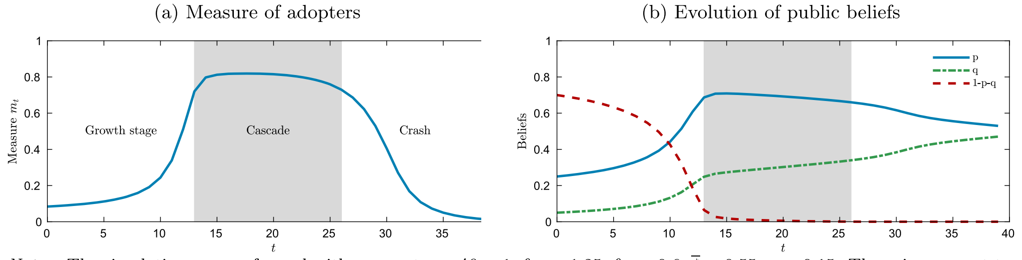

Notes: The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.55$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 3$. The shaded area corresponds to the informational cascade period, loosely defined as periods when $m_t \geq 70\%$ in which learning is markedly slower, given our calibration.

Fig. 4. Endogenous boom and bust with information cascade.

too informative, to prevent agents from learning the truth too quickly, and (iii) the prior q_0 on the false-positive state $\left(\theta_L, \overline{\xi}\right)$ needs to be sufficiently small relative to the true positive $(\theta_H, 0)$ for agents to initially attribute most of the rising adoption pattern to the true positive state.

Fig. 3 presents the evolution of an economy in which the prior about θ is such that the measure of initial adopters is low. When the economy starts in period $t = 0$, the measure of adopters (panel a) is small but higher than agents expected. Seeing a surprisingly high adoption rate, agents understand that it is unlikely to come from the bad state and they reduce the probability assigned to it (dashed curve in panel b). Agents also understand that the high adoption rate could arise from either the good-technology or the false-positive states. As a result, they revise upward their probability assessments of both states (p_t and q_t rise). Importantly, however, given that agents start with a low prior on the false-positive state, the observed high level of adoption is mostly attributed to the good-technology state, so the rise in p_t dominates their expectation. Consequently, agents become more optimistic overall, adoption continues to grow, and the rising adoption pattern, in turn, leads to further upward revisions in expectations, seemingly confirming the assessment that the economy is in the good state. We refer to this first stage of the cycle, characterized by the joint increase of adoption rates and beliefs (p_t, q_t), as the *growth stage*.

Being rational, agents do understand the possibility that they may be mistaken and keep track of the probability of the false-positive state q_t in the background, which also rises throughout the growth stage. Since signals are unbiased along the impulse response path, the belief q_t rises in fact faster than p_t despite starting from a lower prior. Therefore, a time comes when q_t is so high that agents become reluctant to adopt the technology and the aggregate adoption rate begins to decline. This is the start of the *crash* stage, which arises at an endogenous date without the need of an exogenous trigger. As adoption reaches a peak of about 35% given our parametrization, the measure of adopters m_t attains the intermediate region depicted in Fig. 2 where it becomes more informative. As a consequence, agents learn the truth faster, adoption rates drop, and the probability p_t starts declining until a belief reversal occurs later when the belief q_t takes over. Note that the truth is always learned in the end because of the strictly positive information flow.

This example shows that the model is able to generate asymmetric cycles. The growth stage is slow due to the low information flow when m_t is close to 0. The crash, on the other hand, is more sudden because it occurs in the region where i) uncertainty between the good state and the false-positive state is high (p_t and q_t are close, so beliefs are more responsive to new information), and ii) the signal m_t is more informative at the peak.

To highlight the dynamics of the model, the simulations presented in Fig. 3 assume that there are no shocks to the exogenous public signal ($u_t = 0$) and no shock from the noise entrepreneurs ($\varepsilon_t = 0$). But these random shocks, by influencing the signals that agents observe, can also play an important role in driving aggregate adoption rates. In Appendix A.4, we explore how they can affect the timing and intensity of the boom-bust cycle.

Information cascades

Whether the growth stage gives way to a sudden collapse or not depends on the parametrization of the model. Fig. 4 depicts an example of a cycle in which the economy grows so rapidly at first that it reaches the low-informativeness region associated with a high m_t. In that case, the economy goes through an information cascade before the crash. This simulation uses the same parametrization as in Fig. 3, but with a higher common noise parameter $\overline{\xi}$ so that the distribution of private signals in the false-positive and the good states are more similar. Adoption rises faster and reaches higher levels than in the previous example. As a result, there comes a time at the end of the growth stage when agents are so optimistic that they herd on adoption and m_t becomes uninformative. The economy thus enters a period akin to an information cascade, as described earlier, where almost all agents adopt the technology due to overly optimistic public beliefs, and in which markets almost cease to provide information. Through this mechanism, the economy may remain stuck for a long period of time with wrong beliefs and excessive adoption. Because the flow of information is never exactly zero, the economy eventually exits the cascade. This event occurs when the belief about the false-positive state q_t reaches a threshold at which a sufficient fraction of agents stop adopting, bringing back the economy to the region where m_t is informative. The crash takes place in a manner similar to the previous example: beliefs converge more quickly to their true values as the flow of information increases.[13,14]

[13] The way the economy exits the cascade is reminiscent of the "wisdom after the fact" mechanism proposed by Caplin and Leahy (1994) and its reinterpretation in Chapter 4 of Chamley (2004).

[14] The model can also generate permanent information cascades. When private signals s_j are bounded to a set $\left[\underline{s}, \overline{s}\right]$ and the exogenous public signal is uninformative ($\sigma_u \to \infty$), the public beliefs might be so optimistic that the adoption threshold s^* reaches $\underline{s}$. In this case, all agents adopt regardless of their private signals. As a result, the endogenous public signal m_t does not reveal any of the dispersed information and is therefore completely uninformative. The economy is then trapped in a constant state of massive adoption, even though the true quality of the technology is bad. The opposite can happen when public beliefs are sufficiently pessimistic. Appendix A.5 shows how a permanent information cascade can arise in this sample economy.

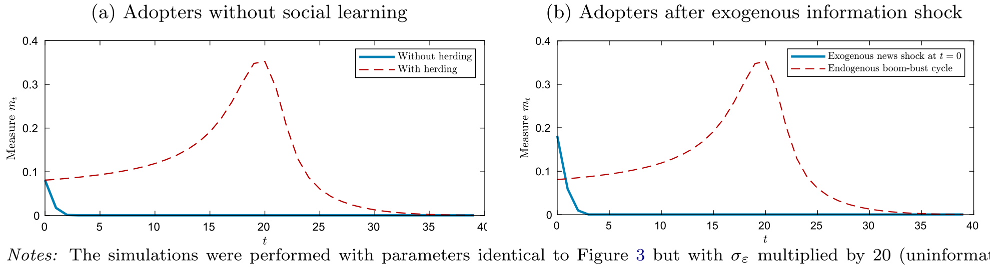

Notes: The simulations were performed with parameters identical to Figure 3 but with σ_ε multiplied by 20 (uninformative endogenous signal). Panel (b) has the same parameters and shows the impact of a shock to u_{-1} so that $p_0 = 0.35$. In both panels the dashed lines represent m_t from (3).

Fig. 5. The role of endogenous learning and the impact of exogenous news shocks.

The importance of herding

To better understand the importance of the social learning channel for the dynamics of the economy, we provide a simulation in which that mechanism is turned off ($\sigma_\varepsilon \to \infty$) so that agents no longer learn by observing m_t and, as a result, no herding takes place. That simulation is presented in panel (a) of Fig. 5. As we can see, without the herding mechanism the economy does not initially grow into a boom, despite being in the false positive state and entrepreneurs having an unusually optimistic prior. In our previous examples, it was the slow diffusion of the information contained in private signals into the public beliefs that led to an economic boom, but since this channel is shut down here, the boom does not happen and the mass of adopters m_t quickly converges to zero.

To further highlight how our endogenous learning mechanism differs from more traditional exogenous noise shocks as in Lorenzoni (2009), we provide another simulation, in panel (b) of Fig. 5, in which the economy is hit by an exogenous u_t shock immediately before $t = 0$ and the social learning channel remains shut down ($\sigma_\varepsilon \to \infty$). Upon impact, this shock leads to an increase in m_t, as expected, but the dynamics it triggers is qualitatively different from that generated by the herding mechanism (dashed curve). As the figure illustrates, adoption peaks immediately then gradually fades out as more information is collected. The propagation is weak and does not lead to the positive feedback loop highlighted in the case of herding, which showed a slow rising pattern of self-reinforcing adoption and optimism.[15]

3.2. Continuum-state case

How general are the phenomena highlighted in the 3-state model? In this section, we discuss under what conditions endogenous boom-and-bust cycles arise in a less restrictive environment.

First, we relax the three-state assumption and return to the specification where ξ can take on a continuum of values. Second, we wish to understand how our two key conditions, (i) $\theta_L + \xi$ close to θ_H and (ii) low q_0, translate to the more general case. To build intuition on this issue, Fig. 6 shows the impulse responses of the economy in the continuous-ξ case assuming that ξ is

[15] Only a specific sequence of increasingly positive then negative exogenous shocks u_t can replicate the type of dynamics observed in Fig. 3. In contrast, in our model, the boom and bust dynamic is the natural outcome of the economic forces at work. The two models also have very different policy implications, as we discuss in Section 3.3.

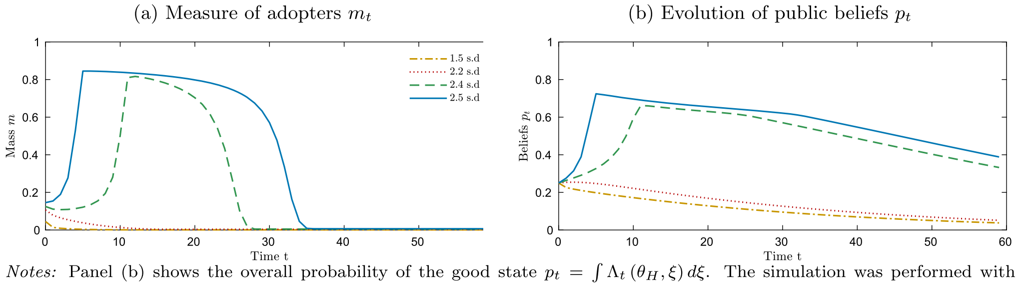

Notes: Panel (b) shows the overall probability of the good state $p_t = \int \Lambda_t(\theta_H, \xi)\, d\xi$. The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian as in Figure 3 with the additional assumption that $\xi \sim \mathcal{N}(0, \sigma_\xi)$ with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 2.5$ and $\sigma_\xi = 0.4$.

Fig. 6. Boom-and-bust cycles in the continuous case.

independent of θ and is normally distributed with mean 0 and standard deviation σ_ξ. As in the previous section, we present the response of the economy in the bad-technology state $\theta = \theta_L$ but we vary the size of the ξ shock. Four shocks of various sizes are represented, with ξ expressed as a multiple of the standard deviation, namely $\xi = k\sigma_\xi$, $k \in \{1.5, 2.2, 2.4, 2.5\}$. The figure shows very distinct behaviors depending on the size of the shock. When the shock is relatively small, $\xi = 1.5\sigma_\xi$ (dash-dotted line), the economy does not experience any herding behavior in which the high initial adoption rate leads to rising optimism. Agents put a sufficiently high likelihood on this ξ draw and are, consequently, able to detect it relatively quickly. Things start to differ as we increase the size of the shock. For an intermediate-sized shock, $\xi = 2.4\sigma_\xi$ (dashed line), the economy begins to experience a boom-bust cycle of the sort described earlier. Because of the low probability of experiencing a shock greater than two standard deviations, agents are initially fooled by the high adoption rates and the economy enters a growth stage with rising optimism and adoption. The growth stage is slow and the crash begins around date $t = 12$, as in Fig. 3. When the size of the shock is larger, $\xi > 2.5\sigma_\xi$ (continuous line), the rise in adoption is so large that the economy goes through an information cascade after experiencing a short growth stage, as in Fig. 4. The economy exits the cascade endogenously at a date which is further delayed as the size of the shock increases.

These simulations show that the dynamics depicted in the examples of Figs. 3 and 4, in the previous section, are not mere curiosities but regular fixtures of the more general model. They also emphasize the importance of *nonlinearities* in governing these dynamics. Indeed, the simulations show that the endogenous boom-and-bust phenomenon occurs whenever the shock to ξ is unusually large, sufficiently so that agents underestimate its likelihood and initially attribute the observation of high adoption rates to the good-technology state.

We now show that there always exists a sufficiently large shock in ξ to trigger a boom-and-bust cycle in beliefs, as long as the exogenous signal coming from the public signal S_t is not too precise.

Proposition 2. *In the Gaussian case, i.e., $F^\xi \sim \mathcal{N}\left(0, \sigma_\xi^2\right)$, $F^s|\theta, \xi \sim \mathcal{N}\left(\theta + \xi, \sigma_s^2\right)$, $F^\varepsilon \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$, $F^u \sim \mathcal{N}\left(0, \sigma_u^2\right)$, for θ and ξ independent and signal S_t sufficiently uninformative (σ_u low), there exists a $\underline{\xi}$ such that all shocks $\xi \geq \underline{\xi}$ generate a boom-and-bust cycle in the impulse response of beliefs p_t to a false-positive shock (θ_L, ξ).*

Note that the above discussion shows a restriction imposed by the theory: because they are rational, agents cannot make systematic mistakes in their assessment of the probability of each state. Hence, boom-bust cycles can only arise for shocks that have a low enough probability of occurring. Our model thus offers a theory of *infrequent* booms-and-busts. Going beyond this limitation may require the introduction of deviations from rationality.

3.3. Welfare

We now turn to the analysis of welfare in this economy. Since entrepreneurs do not internalize that their adoption decisions affect the release of public information, the equilibrium is in general not efficient and policy interventions can be beneficial. To show this formally, we introduce a social planner that maximizes aggregate welfare under limited information. Specifically, we assume that the planner only observes signals that are publicly available and cannot rely on the private information of the entrepreneurs when making decisions.[16]

We follow Angeletos and Pavan (2007) in assuming that the planner seeks to maximize the sum of the entrepreneurs' expected utility, where the expectation is computed according to the agents' private beliefs. In each period, the planner picks an adoption threshold p_t^* such that agents with beliefs $p_{jt} \geq p_t^*$ adopt the technology. Written in recursive form, the problem of the social planner is

$$V(\mathcal{I}) = \max_{p^*} E_{\theta,\xi}\left[\int_{p_j \geq p^*} E\left[A^n \mid \mathcal{I}_j\right] dF_{\theta+\xi}^{p_j}\left(p_j\right) + A^o F_{\theta+\xi}^{p_j}\left(p^*\right) \mid \mathcal{I}\right] + \beta E_{\theta,\xi}\left[V\left(\mathcal{I}'\left(p^*\right)\right) \mid \mathcal{I}\right], \tag{11}$$

where $\mathcal{I}'$ is public information next period, which evolves according to the law of motion (10), and where $F_{\theta+\xi}^{p_j}\left(p_j\right)$ is the CDF of the agents' subjective probability that $\theta = \theta_H$ when the true state of the world is $\theta + \xi$. The expectation $E_{\theta,\xi}$ is then taking over these states using the public beliefs.

The first term in (11) captures the current-period returns from letting agents with private beliefs above p^* adopt the technology. To compute that term, the planner first uses the public beliefs $\mathcal{I}$ to evaluate the likelihood of being in a given state $\theta + \xi$. Since the planner knows the structure of the economy, it can then reconstruct the distribution $F_{\theta+\xi}^{p_j}\left(p_j\right)$ of private beliefs in that state, which is needed to compute the mass of entrepreneurs above and below p^*. The second term in (11), the continuation value, captures the impact of a given adoption threshold p^* on the future public information. It is this term that creates a gap between the equilibrium and the efficient allocation. In the competitive equilibrium, individual agents are atomistic, so they disregard the impact that their actions have on the release of public information. The planner, on the other hand, understands that by changing the cutoff p^*, the mass of adopters also changes, which affects the informativeness of public signals.

The first-order condition of the planner with respect to p^* can be written as

$$E_{\theta,\xi}\left[\lambda\left(p^*\theta_H + \left(1-p^*\right)\theta_L - A^o\right) f_{\theta+\xi}^{p}\left(p^*\right) \mid \mathcal{I}\right] = \beta \frac{\partial E_{\theta,\xi}\left[V\left(\mathcal{I}'\right) \mid \mathcal{I}\right]}{\partial p^*}. \tag{12}$$

[16] We impose these restrictions so that the problem of the planner is not trivial, and that it resembles that of a government trying to design policy under uncertainty about the true value of a new technology.

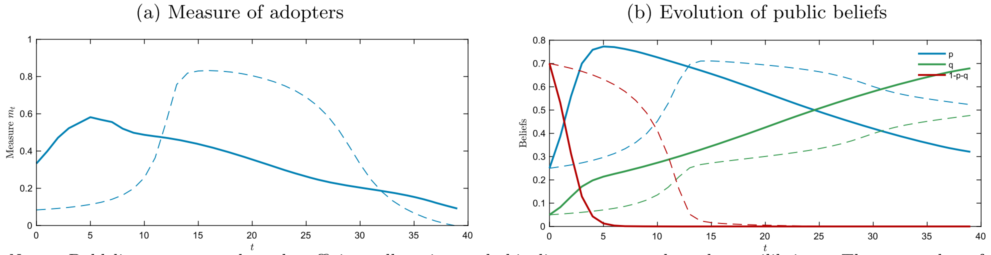

Notes: Bold lines correspond to the efficient allocation and thin lines correspond to the equilibrium. The true value of the fundamental is $(\theta, \xi) = \left(\theta_L, \overline{\xi}\right)$, the false-positive state. The simulation was performed with the same parameters as for Figure 4 and with $\beta = 0.9$ and $\lambda = 0.9$.

Fig. 7. Endogenous boom-bust cycles in the efficient allocation.

The left-hand side of this equation reflects the expected cost of increasing the threshold p^* at the margin. If the true state is $\theta + \xi$, increasing p^* slightly pushes a mass $f^p_{\theta+\xi}(p^*)$ of agents away from adopting, each of which loses $\lambda\left(p^*\theta_H + (1 - p^*)\,\theta_L - A^o\right)$ in expected returns. The planner takes the expectation of these losses over all the states $\theta + \xi$. The right-hand side of the equation reflects the impact of increasing p^* on the flow of public information that is released at the end of the period. By changing p^*, the planner can, for instance, increase the gap between the expected realizations of m in different states of the world. When it does so, m becomes more informative as the signal-to-noise ratio increases. Notice that when $\beta = 0$ the first-order condition (12) collapses to the equilibrium cut-off rule (6), such that the efficient allocation coincides with the equilibrium.

We describe in Appendix A.6 how the efficient allocation can be implemented as an equilibrium through an adoption tax τ^*. That tax in general varies with public beliefs and pushes agents to internalize the social benefits of their adoption decisions on the release information.

Example: Efficiency in the 3-state model

To further explore the role played by inefficiencies in the model, we now go back to our earlier example of Fig. 4, and compare the evolution of the equilibrium and the efficient allocations. The results are presented in Fig. 7 where, as before, the economy is in the false-positive state. We see from Panel (a) that, in the equilibrium (dashed lines), agents are initially cautious and only about 10% of them adopt the new technology. To make the public signal more informative, the planner pushes more agents to adopt (solid lines). It follows that the public beliefs move more rapidly in the efficient allocation (Panel b), and agents quickly learn that the bad-technology can be ruled out. In later periods, starting from about $t = 11$, the situation changes as optimistic signals about the economy accumulate. The public beliefs are so positive that private agents would neglect their private information to adopt the technology massively. The planner instead is more cautious and the efficient allocation features less adoption than in the equilibrium. The planner leans-against-the-wind so that more private information is transmitted into public signals, and that beliefs can move more rapidly. Overall, we see that in the planner's allocation there is a brief boom followed by a steady decline that begins at $t = 5$. In the equilibrium, however, the economy enters an information cascade and aggregate adoption rates remain very high for an extended period. In Appendix A.6 we describe the optimal tax that would need to be implemented in this economy so that the equilibrium coincides with the efficient allocation.

4. A business cycle model with herding

We now embed the mechanism from the simple model into a business cycle framework. Our objective is threefold. First, we want to examine the robustness of the mechanism in a more realistic environment that involves more moving parts (e.g., prices and constraints). Second, we want to investigate under what conditions the hump-shaped evolution of beliefs produced by our simple learning model may lead to a contraction deep enough to go below the trend. Finally, a more realistic setup is required to explore the quantitative implications of the theory, which we do in the next section.

Foreword

A key lesson from the news (or noise)-driven business cycle literature is that standard models have difficulty producing realistic business cycles out of fluctuations in beliefs. In particular they are unable to generate a general macroeconomic expansion, with positive comovements in macro aggregates, during the optimistic phase of the cycle, followed by a recession when beliefs are reversed.

To circumvent this difficulty, we propose a model that features nominal rigidities and two types of capital. Following Lorenzoni (2009), nominal rigidities ensure that aggregate output and hours may increase in response to a surge in demand driven by optimistic expectations. Under a sufficiently accommodative monetary policy, in line with standard estimates, the economy is demand-driven and a muted response of interest rates helps sustain the expansion in demand caused by a positive wealth effect.

We also include two types of capital, a new-technology-specific capital (e.g., IT capital) and a traditional form of capital, in the model to achieve two objectives. First, under the assumption that the new technology is intensive in IT capital, investment must increase for agents to benefit from the technological innovation, ensuring positive comovements between consumption and investment along the booming phase of the cycle. Second, the two types of capital also allow the model to generate a decline below trend during the bust through a *misallocation* channel: when people realize that they overestimated the quality of the technology, the IT capital stock loses its value and the economy must work its way back up to trend by reinvesting in traditional capital.

4.1. Household

Our business cycle framework builds on the core entrepreneurial economy presented in Section 2. In addition to the entrepreneurs who face a technology adoption choice, the economy is composed of i) a representative household, ii) retailers, who are the only agents facing price rigidities, and iii) a monetary authority. The household lives forever, consumes, supplies labor and is the owner of all the firms and capital stocks in the economy. The preferences of the household are given by

$$E\left[\sum \beta^t \log\left(C_t - \frac{L_t^{1+\psi}}{1+\psi}\right)\right], \quad \psi \geq 0,$$

where C_t is the consumption of the final good and L_t is labor. The household can save in a risk-free one-period nominal bond B_t in zero net supply and in two different forms of capital: a traditional type (T) in quantity K_t^T and IT capital in quantity K_t^{IT}. The household is subject to the real budget constraint

$$C_t + \sum_{i=T,IT} I_t^i + \frac{B_t}{P_t} = w_t L_t + \sum_{i=T,IT} z_t^i K_t^i + \frac{1+R_{t-1}}{1+\pi_t}\frac{B_{t-1}}{P_{t-1}} + \Pi_t,$$

where I_t^i, $i = T, IT$, is the investment in each capital type, z_t^i is the corresponding real rental rate, w_t is the real wage, Π_t is total profits, R_{t-1} is the nominal interest rate on government debt issued at date $t-1$, P_t is the nominal price level and $1+\pi_t = P_t/P_{t-1}$ is the inflation rate. As usual, the law of motion for each type of capital, $i = T$ or IT, is given by

$$K_{t+1}^i = (1-\delta)\, K_t^i + I_t^i,$$

where δ is the depreciation rate.

4.2. *Technology*

There are four sectors: i) an entrepreneur sector, ii) a wholesale sector, iii) a retail sector and iv) a final good sector. The most important one, the entrepreneur sector, is the analog of the simple model from Section 2.

Entrepreneur sector

As in the simple model, the entrepreneurial sector is populated by a unit continuum of entrepreneurs indexed by $j \in [0, 1]$ who face a choice between an *old* versus a *new* technology. As before, the old technology is characterized by a known constant productivity A^o and the new technology by an unknown stochastic return A_t^n. In contrast with the earlier model, however, we now assume that these productivities can be combined with capital and labor in a production function. The old production technology is Cobb-Douglas in some capital bundle K_{jt}^o, to be described shortly, and labor L_{jt}^o,

$$Y_{jt}^o = A^o \left(K_{jt}^o\right)^\alpha \left(L_{jt}^o\right)^{1-\alpha}, \; 0 \le \alpha \le 1.$$

Unexpectedly, at date 0, the new technology becomes available with production function

$$Y_{jt}^n = A_t^n \left(K_{jt}^n\right)^\alpha \left(L_{jt}^n\right)^{1-\alpha}.$$

As before, the TFP of the new technology A_t^n is initially as productive as the old one, $A_t^n = A^o$, before it matures with fixed probability $\lambda > 0$. After maturation, the true nature of the technology is revealed and A_t^n is then characterized by a constant fundamental $\theta \in \{\theta_H, \theta_L\}$, $\theta_H > A^o > \theta_L$.

In addition to differing in TFP, the two technologies differ in the capital bundle they use as input. The capital bundle used by each technology $i = o, n$ is given by

$$K_{jt}^i = \kappa_i \left(\omega_i \left(K_{it}^{IT}\right)^{\frac{\zeta-1}{\zeta}} + (1-\omega_i)\left(K_{it}^T\right)^{\frac{\zeta-1}{\zeta}}\right)^{\frac{\zeta}{\zeta-1}}, \; \zeta > 0, \tag{13}$$

where $\kappa_i = \left(\omega_i^\zeta + (1-\omega_i)^\zeta\right)^{-\frac{1}{\zeta-1}}$ and with the assumption that the intensity in IT capital is greater for the new than for the old technology, $0 \le \omega_o < \omega_n \le 1$.[17] We denote by z_t^i, $i = o, n$, the rental price of each bundle.

[17] The value of the parameter κ_i is set such that permanent changes in the measure of new technology adopters m_t have no effect on steady-state output for equal productivities.

After date $t \geq 0$, entrepreneurs face a technology choice problem. As before, we assume that a fraction $0 \leq \mu \leq 1$ of entrepreneurs are "noise entrepreneurs", and that a fraction ε_t of them adopt the new technology, where ε_t is i.i.d., distributed according to a CDF F^ε with support [0, 1]. The remaining $1-\mu$ entrepreneurs are rational and choose the best of the two technologies, based on public and private information:

$$i_{jt} = \underset{i_{jt} \in \{0,1\}}{\operatorname{argmax}}\, i_{jt} E\left[\Pi_t^n \mid \mathcal{I}_{jt}\right] + \left(1 - i_{jt}\right) E\left[\Pi_t^o \mid \mathcal{I}_{jt}\right],$$

where Π_t^i, $i = o, n$, are the profits from using technology i, $\mathcal{I}_{jt}$ is the information set of entrepreneur j at time t, and i_{jt} is a dummy capturing the technology adoption decision. Finally, we assume that entrepreneurs are monopolistic producers of differentiated varieties that sell their output to the wholesale sector.

Wholesale sector

The wholesale, retail and final good sectors play no major role in the model other than separating price rigidities from the technology choice problem of the entrepreneurs.

The wholesale sector is modeled as a representative firm which produces a wholesale good with CES technology

$$Y_t^w = \left(\int_0^1 Y_{jt}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \ \sigma \geq 0, \tag{14}$$

where Y_{jt} is the quantity of inputs it purchases from the monopolistic entrepreneurs. The wholesale sector is perfectly competitive, giving rise to the demand schedule, $Y_{jt} = \left(P_{jt}/P_t^w\right)^{-\sigma} Y_t^w$, where $P_t^w = \left(\int_0^1 P_{jt}^{1-\sigma} dj\right)^{\frac{1}{1-\sigma}}$ is the price of the wholesale good and P_{jt} the price of each differentiated entrepreneur good.

Retail sector

The retail sector is composed of a unit continuum of monopolistic producers who buy the wholesale good at P_t^w and costlessly differentiate it using a one-to-one technology. Retail sector firms are the only ones to face price rigidities. We assume that they face Calvo-style frictions: firms can only reset their price with probability $1-\chi$, leading to a standard Phillips curve.

Final good sector

The final good sector, similar to the wholesale sector, is modeled as a representative firm that operates under perfect competition and produces the final good, used for consumption and investment, using inputs from the retail sector. It uses the CES technology,

$$Y_t = \left(\int_0^1 \left(Y_{jt}^r\right)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}},$$

where Y_{jt}^r is the quantity purchased from each retail firm j and σ is the same elasticity of substitution as in (14).

4.3. Monetary authority

To close the model, we need to specify the policy followed by the monetary authority. As is common in the literature, we assume that the central bank follows a Taylor rule,

$$\frac{1+R_t}{1+\overline{R}} = \left(\frac{1+\pi_t}{1+\overline{\pi}}\right)^{\phi_\pi} \left(\frac{Y_t}{\overline{Y}}\right)^{\phi_y}, \tag{15}$$

where $\overline{R}$, $\overline{\pi}$ and $\overline{Y}$ correspond to the values of, respectively, the target nominal interest rate, inflation and output, which we define later.

4.4. Information

The information structure for entrepreneurs mimics the one in the simple model of Section 2. As before, the true technology θ and the common noise ξ are drawn once-and-for-all at date 0. The ex-ante probability that $\theta = \theta_H$ is denoted by p_0. Entrepreneurs receive a private signal s_j drawn from $F_{\theta+\xi}$ at date 0. In addition, entrepreneurs and all other agents in the economy (household, central bank, retailers, etc.) collect public information over time through the observation of an exogenous signal S_t and the measure of entrepreneurs $m_t = \int_0^{1-\mu} i_{jt} dj + \mu\varepsilon_t$. Importantly, because the productivity of the new technology is identical to A^o until maturation, there is no other source of information in the economy. In equilibrium, prices and aggregate quantities are solely functions of m_t and of the public information up to time t. As a result, prices and quantities provide no other information than is already contained in m_t. The updating rules for beliefs Λ_t and Λ_{jt} are identical to those in the simple model and given by equations (3), (9) and (10).

4.5. Timing

Before date 0, the economy is in a deterministic, no-inflation initial steady state that corresponds to the economy before the introduction of the new technology. At date 0, the new technology fundamental θ, the common noise component ξ and the private signals s_j are drawn once-and-for-all. For all date $t \geq 0$,

1. Entrepreneurs choose whether to adopt the new technology based on the capital stocks $\left(K_t^{IT}, K_t^T\right)$, the relative rental rates z_t^i, $i = T, IT$, and their information set $\left(\Lambda_t, s_j\right)$ (Stage A),
2. The measure of technology adopters m_t is realized,
3. The new technology matures with probability λ,
4. Simultaneously (Stage B),

    (a) All agents observe $\{m_t, S_t\}$ and update their information,
    (b) The household chooses consumption, investment and labor supply,
    (c) Production takes place,
    (d) The monetary authority sets the policy rate,
    (e) Markets clear.

The notation "Stage A" and "Stage B" is used in Appendix B.1 to identify when decisions are made in the full equilibrium definition.

4.6. Investment decision

The technology adoption decision is more complicated than in the simple model because of the presence of general equilibrium effects. When choosing whether to use the new technology, agents have to forecast the profits from either technology. Profits in equilibrium depend not only on productivity but also on the level of demand from wholesalers Y_t^w, prices and the real marginal costs $mc_t^i = \frac{1}{A_{jt}} \left(\frac{z_t^i}{\alpha}\right)^{\alpha} \left(\frac{w_t}{1-\alpha}\right)^{1-\alpha}$ from using each technology $i = o, n$:

$$\Pi_t^i = \left(P_{jt} - P_t mc_t^i\right) Y_{jt}. \tag{16}$$

Solving the model by linearizing the equations of the DSGE model, entrepreneur j ultimately chooses to invest if and only if

$$E\left[\hat{A}_t^n - \alpha \hat{z}_t^n \mid \mathcal{I}_{jt}\right] \geq E\Big[\underbrace{\hat{A}^o}_{=0} - \alpha \hat{z}_t^o \mid \mathcal{I}_{jt}\Big], \tag{17}$$

where the hatted variables are log-deviations from a moving steady state that we define in Section B.3 of the appendix and z_t^i, $i = n, o$, are the rental rates on the capital bundles (13). As equation (17) demonstrates, entrepreneurs not only have to forecast the technology A_t^n but also factor prices, as they are now competing for the same inputs.

This concludes the exposition of the general business cycle model. Additional details with a full definition of the equilibrium, all model equations and a discussion of the resolution method can be found in Appendix B.

5. Quantitative illustration

We now turn to a quantitative illustration of our general macroeconomic model. We calibrate the model to a specific episode in US history and examine its ability to explain broad features of the data.

5.1. Calibration

As we argued before, our model offers a theory of infrequent endogenous booms-and-busts. For that reason, we do not expect to explain general business cycle patterns in the absence of other shocks, but rather to provide a narrative for certain episodes. We thus focus our calibration exercise on a particular episode in recent US history that best fits the description of a technology-driven boom and bust cycle: the late 1990's dot-com bubble. We map the new technology in our model to the introduction of IT technologies in the 1990s and we focus more specifically on the late part of the cycle which covers the period that preceded the stock market collapse in the NASDAQ composite index starting from a trough in 1995Q4 to the crash in 2001Q1.

Our goal with this exercise is not to show that the mechanism can precisely replicate the behavior of the economy during the dot-com period. Rather, we want to determine whether a reasonable calibration of the model is able to generate boom-bust cycles that are similar in terms of magnitude and comovements to what we see in the data. Since we focus on a single historical episode, there is limited data to pin down certain parameters with confidence. For those, we rely on the best data available and provide robustness tests in Appendix B to show that the mechanism does not hinge on specific parameter values. We are, however, careful to properly

Table 1
Standard parameters.

Parameter	Value	Target
α	0.36	Labor share
β	0.99	4% annual interest rate
ψ	2	Frisch elasticity of labor supply (Chetty et al., 2011)
χ	0.75	1 year price duration
σ	10	Markups of about 11%
ϕ_y	0.125	Clarida et al. (2000)
ϕ_π	1.5	Clarida et al. (2000)
ζ	1.71	Elasticity between types of capital (Boddy and Gort, 1971)

discipline moments that are key for the mechanism, such as the dispersion of private beliefs, as we explain in more details below.

The model is solved at a quarterly frequency. Table 1 lists a first set of standard parameters that we take from the literature. The labor intensity α is set to target a standard labor share of 36%. The discount factor β matches an annual real interest rate of about 4%. The household's preference over consumption is logarithmic and the Frisch elasticity is set to 2, within the range of standard macro-level estimates (Chetty et al., 2011). The Calvo price-setting parameter χ yields a standard average price duration of 1 year (Basu and Bundick, 2017). The elasticity of substitution between varieties σ is set to 10 to match an average markup of 11%. The Taylor rule parameters $\left(\phi_y, \phi_\pi\right)$ are within the estimates of Clarida et al. (2000). The target inflation rate $(\overline{\pi})$ is set to 0 and that for output $\left(\overline{Y}\right)$ and the interest rate $\left(\overline{R}\right)$ are set to their respective values in the zero-inflation steady-state before the arrival of the technology. Finally, we pick the elasticity ζ between the different types of capital within the firm from early estimates by Boddy and Gort (1971).

Table 2 lists the more important parameters that attempt to match features of the dot-com bubble. We set the IT-capital shares ω_i, $i = o, n$, to match a share of IT capital of 3% before the introduction of the new technology in 1991 and 14% in 2007 (Strauss and Samkharadze, 2011). The probability of maturation for the new technology λ is set to $1/22$ to match an average waiting time of 22 quarters, corresponding to the length of our period of interest: 1995Q4-2001Q1. We now turn to the technology parameters. A^o is normalized to 1. We use the Survey of Professional Forecaster (SPF) mean real GDP growth forecast over the current quarter to discipline θ_H and θ_L. Under the assumption that factors are fixed in the short run, this identifies changes in the productivity parameter θ. The highest forecast for growth was 4.19% in 2000Q2 in annualized terms. Correcting for a mean growth trend in GDP of 2.4% over 1991-1998, this yields $\theta_H = 1.099$. Similarly, targeting the lowest growth forecast of 0.80% in 2001Q1, we obtain an estimate of $\theta_L = 0.912$.[18]

For the common noise, we adopt the same structure as in Section 3.1 and assume that the pair (θ, ξ) can only take the three values

$$(\theta, \xi) \in \left\{(\theta_L, 0), (\theta_H, 0), \left(\theta_L, \overline{\xi}\right)\right\} \text{ with } \theta_L < \theta_L + \overline{\xi} < \theta_H,$$

[18] The point estimate $\theta_L < A^o$ is not to be interpreted literally as implying that IT technologies are unproductive *per se*, but merely that the market was not fully ripe in 2000 for an economy-wide adoption due to temporary disruption costs (habits, processes) and because other complementary investments (e.g., infrastructure, workforce quality), not modeled here, had not taken place yet. See Appendix B.4 for more details about how we calibrate θ_L and θ_H.

Table 2
Dot-Com episode related parameters.

Parameter	Value	Target
ω_o	0.11	Share of IT capital 1991 (3%)
ω_n	0.26	Share of IT capital 2007 (14%)
λ	1/22	Duration of NASDAQ boom-bust 1995Q4-2001Q1
θ_H	1.099	SPF highest growth forecast over 1998-2001
θ_L	0.912	SPF lowest growth forecast over 1998-2001
s_j	$\mathcal{N}(\theta+\xi, 0.156)$	SPF avg. dispersion in forecasts over 1998-2001
ε	Beta(2, 2)	Non-uniform distribution over [0, 1]
μ	15%	Fraction of noise traders
p_0	0.20	Prior on the "good technology" state
q_0	0.15	Prior on the "false positive" state
$\overline{\xi}$	$0.95(\theta_h-\theta_l)$	See text
σ^u	$3\times\sigma^s$	See text

where $(\theta_L, 0)$ is the *bad-technology* state, $(\theta_H, 0)$ is the *good-technology* state and $\left(\theta_L, \overline{\xi}\right)$ is the *false-positive* state. As before, we let p_t and q_t denote the public beliefs about the good-technology and false-positive states. The distribution of private signals is assumed to be Gaussian, centered on $\theta+\xi$ with standard deviation σ_s. To set the dispersion σ_s, we target the average dispersion of growth forecasts in the SPF over 1995Q4-2001Q1, which yields $\sigma_s = 0.156$. Finally, we must assume a distribution for the fraction of noise traders that adopt the new technology with support over [0, 1]. We choose a beta distribution with parameters (2,2).[19]

Five parameters remain to calibrate for which there does not exist widely accepted estimates or natural targets. The first one is the fraction of noise entrepreneurs μ, which controls the informativeness of the social learning channel. While some estimates exist in the literature regarding the informativeness of markets (see for instance David et al., 2016), these estimates do not cover social learning about new technologies. We conduct sensitivity analysis on this parameter but start with a benchmark value of $\mu = 15\%$. We must also specify the priors p_0 and q_0 that agents associate with the states of the world "good technology" and "false positive". There is no obvious moment that we can target for these parameters given that we focus on one historical episode. We start with a benchmark parametrization that assigns relatively low values to those parameters and do some robustness analysis in the appendix. Our benchmark parametrization attempts to capture the idea that if new technologies can be invented frequently, only few of them lead to deep economic transformations likes the ones considered in this paper. This suggests a small value for p_0. Similarly, the false-positive signal ξ cannot happen too frequently, otherwise agents would distrust their private signals and the boom-bust cycles would never arise. As a benchmark, we therefore pick $p_0 = 0.2$ and $q_0 = 0.15$ but we show in Appendix B that boom-bust episodes are robust to some variation in these values. We must also set a value for the common noise term $\overline{\xi}$. Again, neither the literature nor the data provides much guidance. Since one of our goals is to evaluate the potential of the model in generating boom-bust cycles, we pick a relatively large shock and set $\overline{\xi} = 0.95\,(\theta_H - \theta_L)$. Finally, we need a value for the exogenous public signal about θ. For the mechanism to operate, we need that signal to be mostly uninformative, otherwise agents know the true value of θ and boom-bust cycles can obviously never arise. We set

[19] The distribution Beta(1, 1) is uniform and produces a flat learning response. As a result, we pick a Beta(2, 2) distribution which is symmetric around its mode of 0.5. This assumption is relatively unimportant.

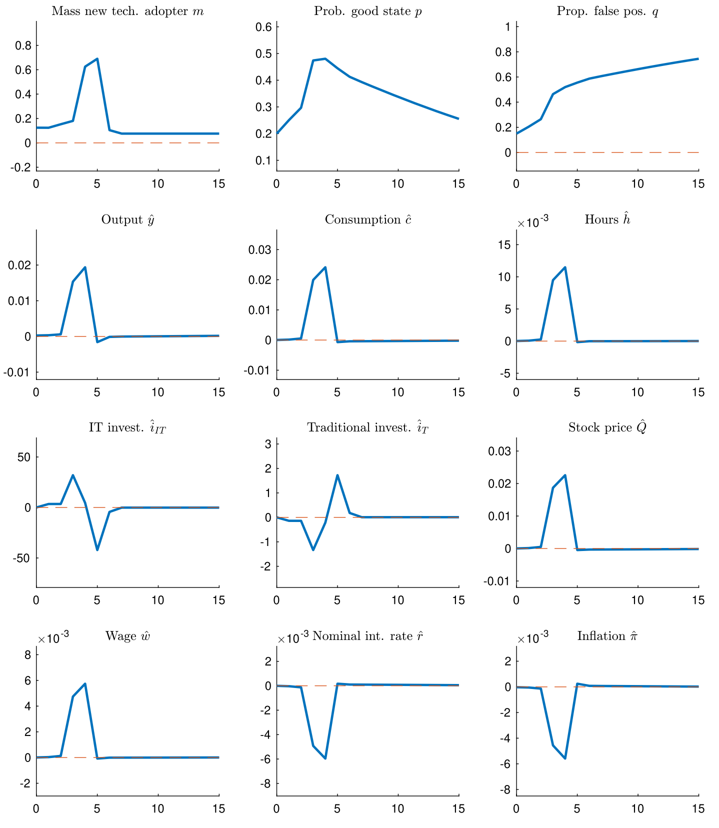

Notes: The impulse responses are reported in log-deviations from the initial non-stochastic steady state. All shocks after the realization of (θ, ξ) are set to 0 ($\varepsilon_t = 0$ and the technology never matures exogenously).

Fig. 8. Impulse response in the false-positive state of the world.

$\sigma_u = 3 \times \sigma_s$ in the benchmark simulations so that a private signal has about the same information as three exogenous public signals. Appendix B conducts sensitivity analysis over all the parameters mentioned in this paragraph. The results are fairly robust.

5.2. Boom-and-bust cycles

Our interpretation of the dot com bubble is that the economy was in the false-positive state, and that investors' optimism about IT technologies triggered the initial boom. As more information became available, agents became pessimistic, and the bust followed.[20] In line with this interpretation, Fig. 8 presents the impulse responses of our model economy in the false-positive

[20] Ofek (2002) looks at the earnings trajectory that would have been required to justify the stock valuation of internet companies in early 2000. He finds that, according to a standard asset pricing model, the earnings growth of internet companies would have had to exceed 40% for a decade to justify their valuations. These large growth rates did not materialize. This analysis is consistent with investors having a very optimistic view about the fundamental that turned out to be mistaken, in line with our false-positive interpretation.

state with all the other shocks are set to zero, that is $\varepsilon_t = 0$ and $u_t = 0$ for all t. We also assume that the technology never matures through the λ shock, so that the bust that we observe is purely endogenous.[21]

In period $t = 0$, a new technology is discovered and entrepreneurs receive encouraging private signals about the true value of that technology. Since the false-positive state is initially deemed unlikely, entrepreneurs begin to adopt the new technology and the economy goes through a growth phase with m_t moving upward. While this growth in m_t is initially consistent with both the "good technology" and the "false-positive" state, there comes a point at which agents start to realize that the data is more consistent with $\theta = \theta_L$, and the likelihood p_t starts to decline. As a result, the mass of entrepreneurs who adopt the new technology collapses around $t = 5$, pushing the economy into a crash.

While the behavior of the mass of adopters m_t is similar to what we have observed in the simplified model of Section 3.1, Fig. 8 shows how this pattern and the evolution of beliefs translate to other macroeconomic variables. As agents become more optimistic after observing people rushing to adopt the new technology, the household anticipates higher productivity growth in the future and higher income, resulting in upward pressure on consumption due to a positive income effect. With expectations of higher productivity from the new technology, the demand for IT capital rises and the household responds by increasing IT investment. The new technology being less intensive in the other form of capital, the demand for traditional capital falls and so does traditional investment. The rise in consumption and investment in IT capital, despite being accompanied by a moderate decline in traditional investment, contribute to an overall rise in aggregate demand. Price rigidities play an important role in turning this surge in demand into a general macroeconomic boom. In a real business cycle model, the rise in aggregate demand should be offset by a sharp rise in the real interest rate. With sticky prices, the interest rate response is muted if the monetary authority is sufficiently accommodative. As a result, aggregate demand remains high. Firms, satisfying demand, respond by raising output and employment. Because of a higher labor demand, wages increase, but inflation remains low because firms anticipate greater productivity and lower marginal costs in the future. As evidenced by the variable $\hat{Q}$, which captures the value of the firms, the economy also experiences a stock market boom along the expansion.

These dynamic effects are reversed when the crash occurs and agents realize that the new technology is actually of low quality. While agents abandon the new technology, a recession occurs, with GDP falling below trend, because agents wake up after having invested too much in IT capital and not enough in traditional capital. This misallocation of resources, combined with a negative income effect, is the essential ingredient that puts downward pressure on aggregate demand and pushes the economy below the trend in the recovery.

A few comments are in order at this point. First, while the model is able to generate a recession with a significant peak-to-trough gap (about 2%), it remains smaller than the one in the data (about 3%). This result seems, however, a feature of belief-driven cycles that our model shares with most of the news/noise-driven business cycle literature. Second, while the model is able to generate a recession with output falling below trend, that effect is somewhat weak in our simula-

[21] An alternative view of the data is that the bust was triggered by the technology maturing. While we cannot completely rule out that story, we note that measures of labor productivity in the technology industry have kept increasing through the bust, which pushes against a scenario in which the IT technology would be revealed to be of low productivity. Also, the literature does not seem to agree, even roughly, on any given event (a potential λ shock) that could have triggered the bust.

tion. Given that the capital stocks can be adjusted rapidly, the misallocation channel responsible for the dip does not lead to a large drop in output. The introduction of debt and bankruptcy in the model would provide another channel through which the crash could result in a deeper recession. Third, we can also compute the frequency at which boom-and-bust cycles arise in our model. While the existing consensus is that such cycles are rare in models with rational agents,[22] our benchmark calibration suggests that boom-bust cycles may arise in our calibrated model at the fairly high frequency of $q_0 = 15\%$ after the introduction of a new technology. We view this number as quite encouraging for the ability of rational herding models in explaining the data.

Overall, the impulse responses of Fig. 8 suggest that it is possible to generate a realistic macroeconomic boom-bust cycle that is entirely driven by the internal forces of the model—without the need of an exogenous shock to trigger the bust after the initial introduction of a new technology.[23]

6. Conclusion

This paper explores whether rational herding can generate endogenous business cycle fluctuations. We propose a novel theory of herding which captures many essential features of more traditional models (Banerjee, 1992; Bikhchandani et al., 1992; Chamley, 2004), while being tractable enough to be embedded into a general equilibrium business cycle framework. We show that the model is able to endogenously generate a boom-and-bust pattern without the need for a particular sequence of shocks. Our model has predictions on the frequency, the timing and the conditions under which such cycles emerge or burst. It can thus be used to analyze the role of stabilization policy, including investment-specific taxes or monetary policy.

We have restricted our attention to technology-driven boom-and-bust cycles, but the implications of the theory go beyond this context and we believe our herding model can be used in other environments to analyze herding behavior following any sort of innovation, be it financial innovations or innovations to the demand for certain types of goods (new products, housing, etc).

Several extensions are worth investigating. First, our current macroeconomic model ignores the role of debt. An interesting extension would be to study how the rising pattern of optimism during the growth stage of the cycle could relax financial constraints and lead to an expansion in credit, triggering a wave of bankruptcies at the time of the crash. Another natural extension would be to consider a financial market application of our herding model and examine, in particular, the role of speculation. We leave these ideas to future research.

Declaration of competing interest

The authors declare that they have no financial or personal interests that relate to the research described in this paper.

Data availability

Data will be made available on request.

[22] Chamley (2004) suggests that a boom-bust cycle arises with a probability of 10^{-6} in the traditional model of herding of Avery and Zemsky (1998).

[23] In Appendix B.5, we investigate the role of lean-against-the-wind monetary policy in this economy and find that it only has a limited impact on boom-bust cycles.

Appendix A. Appendix of Section 2

A.1. Equations for the three-state model

This section provides the specific model equations that characterize beliefs in the three-state model. Equation (3) that builds private beliefs from the public ones becomes

$$p_{jt} = p_j\left(p_t, q_t, s_j\right) = \frac{p_t f^s_{\theta_H}\left(s_j\right)}{p_t f^s_{\theta_H}\left(s_j\right) + q_t f^s_{\theta_L+\overline{\xi}}\left(s_j\right) + \left(1-p_t-q_t\right) f^s_{\theta_L}\left(s_j\right)}, \tag{18}$$

$$q_{jt} = q_j\left(p_t, q_t, s_j\right) = \frac{q_t f^s_{\theta_L+\overline{\xi}}\left(s_j\right)}{p_t f^s_{\theta_H}\left(s_j\right) + q_t f^s_{\theta_L+\overline{\xi}}\left(s_j\right) + \left(1-p_t-q_t\right) f^s_{\theta_L}\left(s_j\right)}.$$

Equation (9) that defines the interim beliefs after observing S_t is simply

$$p_{t|R_t} = \frac{p_t f^u\left(S_t-\theta_H\right)}{p_t f^u\left(S_t-\theta_H\right) + \left(1-p_t\right) f^u\left(S_t-\theta_L\right)},$$

$$q_{t|R_t} = \frac{q_t f^u\left(R_t-\theta_L\right)}{p_t f^u\left(S_t-\theta_H\right) + \left(1-p_t\right) f^u\left(S_t-\theta_L\right)}.$$

Finally, in the three state model, the optimal adoption strategy characterized by Equation (10) that defines the law of motion of beliefs after observing m_t becomes

$$p_{t+1} = \frac{p_{t|R_t} f^{\varepsilon}\left(\left(m_t-(1-\mu)\overline{F}^s_{\theta_H}\left(s^*\left(p_t,q_t\right)\right)\right)/\mu\right)}{denom_t},$$

$$q_{t+1} = \frac{q_{t|R_t} f^{\varepsilon}\left(\left(m_t-(1-\mu)\overline{F}^s_{\theta_L+\overline{\xi}}\left(s^*\left(p_t,q_t\right)\right)\right)/\mu\right)}{denom_t},$$

where

$$\begin{aligned} denom_t = {} & p_{t|R_t} f^{\varepsilon}\left(\left(m_t-(1-\mu)\overline{F}^s_{\theta_H}\left(s^*\left(p_t,q_t\right)\right)\right)/\mu\right) \\ & + q_{t|R_t} f^{\varepsilon}\left(\left(m_t-(1-\mu)\overline{F}^s_{\theta_L+\overline{\xi}}\left(s^*\left(p_t,q_t\right)\right)\right)/\mu\right) \\ & + \left(1-p_{t|R_t}-q_{t|R_t}\right) f^{\varepsilon}\left(\left(m_t-(1-\mu)\overline{F}^s_{\theta_L}\left(s^*\left(p_t,q_t\right)\right)\right)/\mu\right). \end{aligned}$$

A.2. Propositions

Proposition 1. *There exists a unique equilibrium.*

Proof. The threshold p^* is uniquely determined by (6). The result is established recursively. Fix the fundamental (θ, ξ) and the realization of the shocks $\{u_0, \varepsilon_0, u_1, \varepsilon_1, \ldots\}$. At any date t, given public beliefs Λ_t, (3) and (5) yield a unique distribution of private beliefs $\left\{\Lambda_{jt}\right\}_{j\in[0,1]}$ and $\left\{p_{jt}\right\}_{j\in[0,1]}$. Given these, under the tie-breaking rule that indifferent agents adopt the technology, there is a unique m^e_t derived from (8) and, therefore a unique m_t from (7). As a result, updating beliefs through (9) and (10) yields unique $\Lambda_{t|R}$ and Λ_{t+1}. We have shown that the updating of public beliefs yields a unique Λ_{t+1} from Λ_t and the realization of shocks $\{u_t, \varepsilon_t\}$. Starting from public beliefs Λ_0, there is therefore a unique equilibrium path $\{\Lambda_0, \Lambda_1, \ldots\}$ for any history of shocks, and all other quantities can be uniquely determined from it. □

Lemma 1. *In the three-state model, for $\theta_L < \theta_L + \overline{\xi} < \theta_H$ and $\left\{F_x^s\right\}$ satisfying the MLRP condition, the optimal adoption strategy in characterized by a cutoff rule in the private signal $s^*(p_t, q_t)$, decreasing in p_t. That is, an agent adopts the technology if and only if $s_j \geq s^*(p_t, q_t)$. The measure of rational adopters is given by*

$$m^e(p_t, q_t, \theta, \xi) = (1-\mu)\,\overline{F}^s_{\theta+\xi}\left(s^*(p_t, q_t)\right).$$

Proof. The proof is straightforward. Under the above conditions, rewrite the individual probability of the good-technology state as

$$p_j\left(p_t, q_t, s_j\right) = \frac{p_t}{p_t + q_t \frac{f^s_{\theta_L+\overline{\xi}}(s_j)}{f^s_{\theta_H}(s_j)} + (1-p_t-q_t)\frac{f^s_{\theta_L}(s_j)}{f^s_{\theta_H}(s_j)}}.$$

Under the assumption of MLRP and $\theta_L < \theta_L + \overline{\xi} < \theta_H$, p_j is clearly increasing in s_j. Hence, for all (p_t, q_t), there exists a cutoff $s^*(p_t, q_t) \in \mathbb{R} \cup \{-\infty, \infty\}$ such that $s_j \geq s^*(p_t, q_t) \Leftrightarrow p_j\left(p_t, q_t, s_j\right) \geq p^*$. Also, because p_j is increasing in p_t, the implicit function theorem ensures that $s^*(p_t, q_t)$ is decreasing in p_t. The measure of rational adopters is thus

$$\begin{aligned} m^e(p_t, q_t, \theta, \xi) &= (1-\mu)\int \mathbb{I}\left(p_j\left(p_t, q_t, s_j\right) \geq p^*\right) f^s_{\theta+\xi}\left(s_j\right) ds_j \\ &= (1-\mu)\,\overline{F}^s_{\theta+\xi}\left(s^*(p_t, q_t)\right). \quad \square \end{aligned}$$

Proposition 2. *In the Gaussian case, i.e., $F^\xi \sim \mathcal{N}\left(0, \sigma_\xi^2\right)$, $F^s|\theta, \xi \sim \mathcal{N}\left(\theta+\xi, \sigma_s^2\right)$, $F^\varepsilon \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$, $F^u \sim \mathcal{N}\left(0, \sigma_u^2\right)$, for θ and ξ independent and signal S_t sufficiently uninformative (σ_u low), there exists a large enough $\underline{\xi}$ such that all shocks $\xi \geq \underline{\xi}$ generate a boom-and-bust cycle in the impulse response of beliefs p_t to a false-positive shock (θ_L, ξ).*

Proof. Our strategy is to show that there exists a sufficiently large $\underline{\xi}$, such that for all shock $\xi \geq \underline{\xi}$ the public beliefs about the good state in date 1, p_1, increases after observing m_0. Since beliefs must converge to the truth in the long-run ($p_t \to 0$), due to the strictly positive flow of information, and the law of large numbers, this guarantees the existence of a boom-and-bust cycle in beliefs. We start under the assumption that S_t is totally uninformative, $\sigma_u = \infty$.

First, we establish that the optimal strategy in the Gaussian case follows a cutoff strategy in s^*. The probability that individual j puts on the good state is given by

$$p_j\left(p_0, s_j\right) = \frac{\int \Lambda_0(\theta_H, \xi)\, f^s_{\theta_H+\xi}\left(s_j\right) d\xi}{\int \Lambda_0(\theta_H, \xi)\, f^s_{\theta_H+\xi}\left(s_j\right) d\xi + \int \Lambda_0(\theta_L, \xi)\, f^s_{\theta_L+\xi}\left(s_j\right) d\xi}.$$

Since ξ is independent from θ, $\Lambda_0(\theta_H, \xi) = p_0 f^\xi(\xi)$ and $\Lambda_0(\theta_L, \xi) = (1-p_0) f^\xi(\xi)$. Notice, then, that $\int f^\xi(\xi) f^s_{\theta+\xi}\left(s_j\right) d\xi$ is the pdf of s_j given θ, which is a normal, $s_j|\theta \sim \mathcal{N}\left(\theta, \sigma_\xi^2 + \sigma_s^2\right)$. Denote ϕ the pdf of a unit normal, we have:

$$p_j\left(p_0, s_j\right) = \frac{1}{1 + \frac{(1-p_0)\int f^\xi(\xi) f^s_{\theta_L+\xi}(s_j) d\xi}{p_0 \int f^\xi(\xi) f^s_{\theta_H+\xi}(s_j) d\xi}} = \frac{1}{1 + \frac{1-p_0}{p_0}\phi\left(\frac{s_j-\theta_L}{\sqrt{\sigma_\xi^2+\sigma_s^2}}\right)/\phi\left(\frac{s_j-\theta_H}{\sqrt{\sigma_\xi^2+\sigma_s^2}}\right)}.$$

Since the Gaussian family satisfies the MLRP property, p_j is increasing in s_j. Hence, the optimal adoption strategy at date 0 takes a cutoff form $\hat{s}_0$.

Under the assumption that S_t is uninformative, the public belief about the good state at the beginning of period 1, p_1, is given by

$$p_1 = \int \Lambda_1(\theta_H, \xi)\, d\xi$$
$$= \frac{\int \Lambda_0(\theta_H, \xi)\, f^{\varepsilon}\left((m_0 - m^e(\Lambda_0, \theta_H, \xi))/\mu\right) d\xi}{\int \Lambda_0(\theta_H, \xi)\, f^{\varepsilon}\left((m_0 - m^e(\Lambda_0, \theta_H, \xi))/\mu\right) d\xi + \int \Lambda_0(\theta_L, \xi)\, f^{\varepsilon}\left((m_0 - m^e(\Lambda_0, \theta_L, \xi))/\mu\right) d\xi}.$$

Using the independence property between θ and ξ and the cutoff property, the above formula can be rewritten as

$$p_1 = \frac{1}{1 + \frac{1-p_0}{p_0} \frac{\int f^{\xi}(\xi) f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_L+\xi}(s_0^*)\right)/\mu\right) d\xi}{\int f^{\xi}(\xi) f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_H+\xi}(s_0^*)\right)/\mu\right) d\xi}}.$$

Denoting ξ_0 the true shock, the impulse response in m_t yields $m_0 = (1-\mu)\overline{F}^s_{\theta_L+\xi_0}(s_0^*)$, which goes to $1-\mu$ as $\xi_0 \to \infty$. Because the MLRP property implies first-order stochastic dominance, we have $\overline{F}^s_{\theta_L+\xi}(s_0^*) < \overline{F}^s_{\theta_H+\xi}(s_0^*)$. Since $f^{\varepsilon}(\varepsilon)$ is decreasing for $\varepsilon \geq 0$, we have

$$f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_L+\xi}(s_0^*)\right)/\mu\right) < f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_H+\xi}(s_0^*)\right)/\mu\right)$$

for all $\xi \leq \theta_L - \theta_H + \xi_0$. Decompose the difference between the denominator and numerator can be written

$$\begin{aligned}
&\int f^{\xi}(\xi) f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_H+\xi}(s_0^*)\right)/\mu\right) d\xi \\
&\quad - \int f^{\xi}(\xi) f^{\varepsilon}\left(\left(m_0 - (1-\mu)\overline{F}^s_{\theta_L+\xi}(s_0^*)\right)/\mu\right) d\xi \\
&\underset{\xi_0\to\infty}{\longrightarrow} \int_{-\infty}^{\infty} f^{\xi}(\xi) \Big[f^{\varepsilon}\left(\left(1-\mu-(1-\mu)\overline{F}^s_{\theta_H+\xi}(s_0^*)\right)/\mu\right) \\
&\quad - f^{\varepsilon}\left(\left(1-\mu-(1-\mu)\overline{F}^s_{\theta_L+\xi}(s_0^*)\right)/\mu\right)\Big] d\xi > 0
\end{aligned}$$

The difference converges to a strictly positive term. Thus, there exists $\underline{\xi}$ such that for all $\xi_0 > \underline{\xi}$

$$\begin{aligned}
&\int f^{\xi}(\xi) f^{\varepsilon}\left(\left((1-\mu)\overline{F}^s_{\theta_L+\xi_0}(s_0^*) - (1-\mu)\overline{F}^s_{\theta_L+\xi}(s_0^*)\right)/\mu\right) d\xi \\
&< \int f^{\xi}(\xi) f^{\varepsilon}\left(\left((1-\mu)\overline{F}^s_{\theta_L+\xi_0}(s_0^*) - (1-\mu)\overline{F}^s_{\theta_H+\xi}(s_0^*)\right)/\mu\right) d\xi
\end{aligned}$$

and $p_1 > p_0$. The shock is large enough for agents to attribute it mostly to the good state, initiating the growth stage of the cycle. By continuity of the belief updating equations in σ_u. There must also exists a sufficiently large σ_u (S_t sufficiently uninformative) for which $p_1 > p_0$ after $\xi_0 \geq \underline{\xi}$. □

A.3. Additional figures

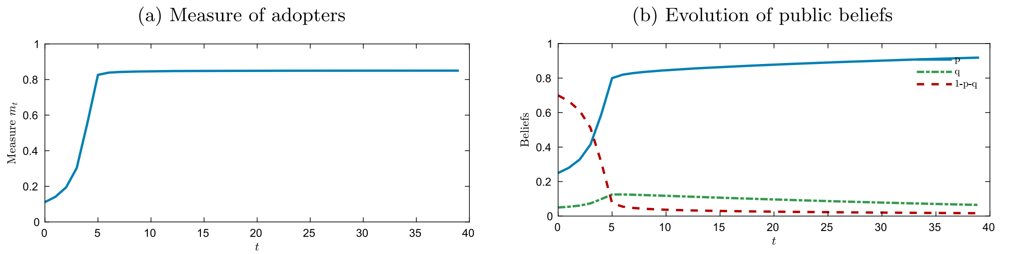

Notes: The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.55$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 3$.

Fig. 9. Impulse response in the case of a true positive.

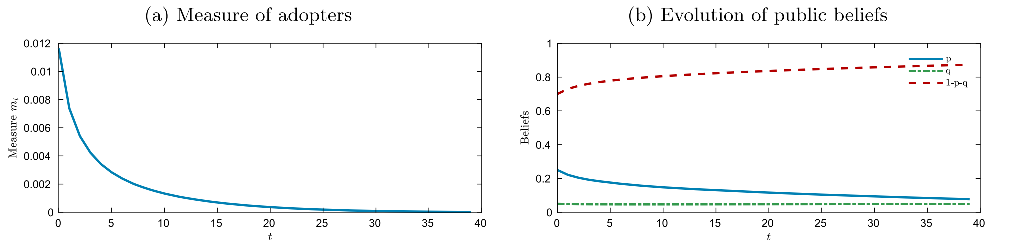

Notes: The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.55$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 3$.

Fig. 10. Impulse response in the case of a true negative.

A.4. The role of random shocks

To highlight the dynamics of the model, the simulations presented in Figs. 3 and 4 assumed that there were no shocks to the public signal S_t ($u_t = 0$) and no shock coming from the noise entrepreneurs ($\varepsilon_t = 0$). But these random shocks, by influencing the signals that agents observe, can also play an important role in driving aggregate adoption. Because single shocks do not have much effect, we illustrate this by conducting a series of simulations using the same economy as in Fig. 3 but in which we fix the shocks to some number $\bar{x}$. As a result, we can see how the economy evolves when agents continuously receive optimistic signals and, inversely, when they get a constant flow of pessimistic news. We plot the results in Fig. 11. In the left panel we set $u_t = 0$ and $\varepsilon_t = \bar{x}$ and vary $\bar{x}$ from -0.005 to 0.005. In the right panel we set $\varepsilon_t = 0$ and $u_t = \bar{x}$ and vary $\bar{x}$ from -0.1 to 0.1. The continuous line represents the simulation with $\bar{x} = 0$, and the dashed lines represent simulations with $\bar{x} \neq 0$. As we can see from the figure, there is quite a bit of dispersion across simulations and shocks to u_t and ε_t can push the economy through very different dynamics. For the more optimistic signals (u_t, ε_t), the economy enters an information cascade and adoption remains high for a sustained period. In contrast, for the more pessimistic signals, adoption slowly declines and we never observe a boom-bust cycle.

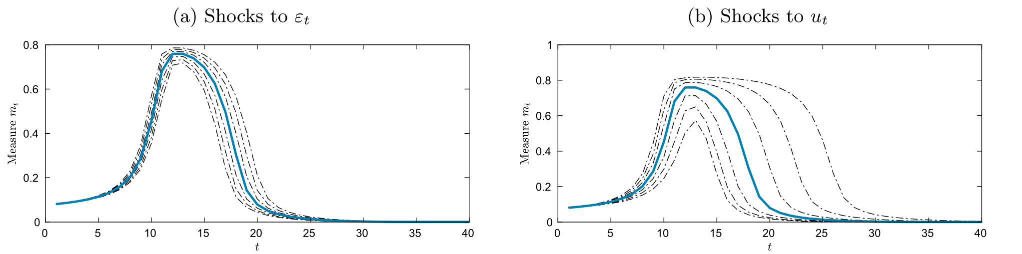

Notes: Simulations performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.53$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = 3$. Left panel: the signals are $u_t = 0$ and $\varepsilon_t = \bar{x}$ and we vary $\bar{x}$ from -0.005 (lowest line) to 0.005 (highest line) in seven increments. Right panel: the signals are $\varepsilon_t = 0$ and $u_t = \bar{x}$ and we vary $\bar{x}$ from -0.1 (lowest line) to 0.1 (highest line) in seven increments.

Fig. 11. The role of u_t and ε_t in driving aggregate adoption.

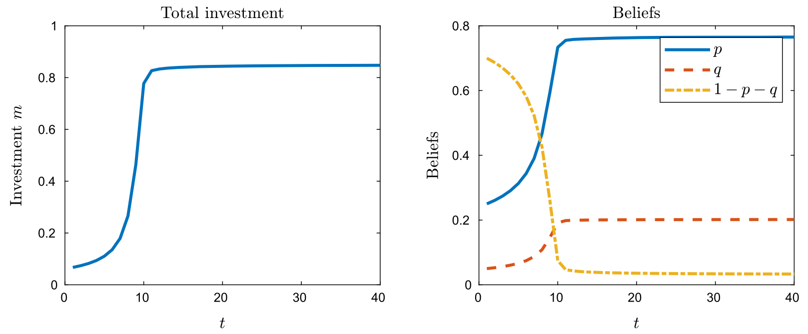

Notes: The simulation was performed with parameters: $A^o = 1$, $\theta_H = 1.25$, $\theta_L = 0.6$, $\overline{\xi} = 0.55$, $\mu = 0.15$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.6$, $\sigma_\varepsilon = 1$, $\sigma_u = \infty$. The private signal distribution is bounded below by zero.

Fig. 12. Impulse response in the false-positive state with permanent information cascades.

Notice also that shocks to μ_t and ε_t, as they both influence the overall level of optimism in the economy, have similar effects on the dynamic of adoption.[24]

A.5. Permanent information cascades

In this appendix, we show that a slight modification of the environment considered in Fig. 4 in the main text can lead to a permanent information case in which no information is provided by the observation of the mass of adopters. We depart from that specification by imposing a lower bound $\underline{s} = 0$ on the distribution of private signal and by making the exogenous public signal uninformative ($\sigma_u = \infty$). Fig. 12 shows the behavior of the economy in the false-positive state under this new parametrization. As we can see, the economy quickly converges to a permanent plateau at $m = 1$. The public beliefs are so optimistic at that point that even the most pessimistic agent ($s_j = 0$) adopts the technology. Since all agents adopt regardless of the true state of the

[24] Some of the curves in Fig. 11 intersect each other. For instance, in the case with $\bar{x} = 0$ agents learn fairly slowly about the true state of the world. In contrast, in the simulation with a slightly higher $\bar{x}$ the behavior of aggregate adoption leads to a high flow of information and agents quickly learn that the true fundament is bad.

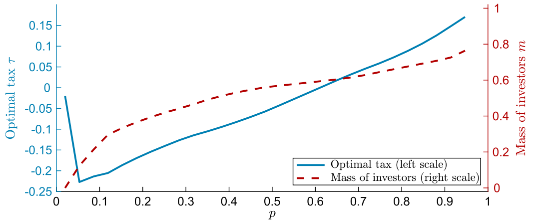

Notes: Mass of adopters and optimal tax as a function of p with the same parameters as Figure 7 and with $\beta = 0.9$ and $\lambda = 0.9$. These curves are generated by fixing $q = 0.01$. m is computed as $(1-\mu)\,\overline{F}_{\theta_H}\left(\hat{s}\left(p,q\right)\right)$.

Fig. 13. Optimal tax as a function of p.

world, observing the mass of adopters provides no information and the economy remains in a permanently elevated state of adoption.

A.6. *Optimal taxation*

To gain further insight into the nature of the model's inefficiencies, it is useful to look at a particular implementation of the efficient allocation using a tax (or subsidy) τ^* that entrepreneurs must pay to adopt the new technology. The next proposition characterizes τ^*.

Proposition 3. *The efficient allocation can be implemented as an equilibrium by an adoption tax*

$$\tau^* = \left(E_{\theta,\xi}\left[f^{p}_{\theta+\xi}\left(p^*\right) \mid \mathcal{I}\right]\right)^{-1} \beta \frac{\partial E_{\theta,\xi}\left[V\left(\mathcal{I}'\right) \mid \mathcal{I}\right]}{\partial p^*}, \tag{19}$$

and a lump-sum transfer to all entrepreneurs.

Proof. We consider a tax τ that agents must pay to adopt the new technology. Under that tax, (6) shows that the marginal adopter p^* is such that $\lambda\left(p^*\theta_H + \left(1-p^*\right)\theta_L\right) + (1-\lambda)\,A^o - \tau = A^o$. Combining with (12) and reorganizing yields (19). □

The optimal tax τ^* balances the distortion in adoption it creates (first term in the product) with the potential benefit on information acquisition (second term).

We plot in Fig. 13 how this optimal tax varies with the public beliefs p in the example of Fig. 7.[25] We see that the tax tends to be negative for low values of p and positive for larger values. As we described in Section 3.1, when agents are pessimistic (low p), few of them adopt the technology and the endogenous public signal does not reveal much information. The planner therefore sets $\tau^* < 0$ to encourage entry and make the observed mass of adopters a more precise signal. The opposite happens when many entrepreneurs adopt (high p). In this case, the planner sets $\tau^* > 0$ to discourage adoption, once again to make the endogenous public signal more informative. The tax for intermediate values of p reflects these information concerns.[26]

In this particular Gaussian case, the tax incentivizes agents to behave against the crowd in what amounts to a *leaning-against-the-wind* pattern: the tax is negative when no agent wants to

25 To draw this plot, we fix q to some arbitrary value $q = 0.01$ and plot m in the good-technology state.

26 The planner begins to phase out the tax as p approaches 0. In this case, the public beliefs are so extreme that there is very little uncertainty about the true state of the world. Since there is not much more to learn, the planner sets the tax close to zero to minimize the distortion in adoption.

adopt the technology, and positive when agents adopt massively.[27] We will see in our quantitative model that these same forces have important consequences for the conduct of monetary policy.

Appendix. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2023.105669.

References

Adam, K., Marcet, A., Beutel, J., 2017. Stock price booms and expected capital gains. Am. Econ. Rev. 107 (8), 2352–2408.
Angeletos, G.-M., La'O, J., 2013. Sentiments. Econometrica 81 (2), 739–779.
Angeletos, G.-M., Pavan, A., 2007. Efficient use of information and social value of information. Econometrica 75 (4), 1103–1142.
Angeletos, G.-M., Lorenzoni, G., Pavan, A., 2022. Wall street and silicon valley: a delicate interaction. Rev. Econ. Stud.
Arif, S., Lee, C.M., 2014. Aggregate investment and investor sentiment. Rev. Financ. Stud. 27 (11), 3241–3279.
Asriyan, V., Fornaro, L., Martin, A., Ventura, J., 2019. Monetary policy for a bubbly world. Working paper.
Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. Am. Econ. Rev., 724–748.
Banerjee, A.V., 1992. A simple model of herd behavior. Q. J. Econ. 107 (3), 797–817.
Barlevy, G., 2014. A leverage-based model of speculative bubbles. J. Econ. Theory 153, 459–505.
Basu, S., Bundick, B., 2017. Uncertainty shocks in a model of effective demand. Econometrica 85 (3), 937–958.
Beaudry, P., Portier, F., 2004. An exploration into Pigou's theory of cycles. J. Monet. Econ. 51 (6), 1183–1216.
Beaudry, P., Portier, F., 2014. News-driven business cycles: insights and challenges. J. Econ. Lit. 52 (4), 993–1074.
Beaudry, P., Galizia, D., Portier, F., 2020. Putting the cycle back into business cycle analysis. Am. Econ. Rev. 110 (1), 1–47.
Benhabib, J., 1992. Cycles and Chaos in Economic Equilibrium. Princeton University Press.
Benhabib, J., Farmer, R.E., 1994. Indeterminacy and increasing returns. J. Econ. Theory 63 (1), 19–41.
Benhabib, J., Schmitt-Grohé, S., Uribe, M., 2002. Chaotic interest-rate rules. Am. Econ. Rev. 92 (2), 72–78.
Benhabib, J., Wang, P., Wen, Y., 2015. Sentiments and aggregate demand fluctuations. Econometrica 83 (2), 549–585.
Benhima, K., 2019. Booms and busts with dispersed information. J. Monet. Econ. 107, 32–47.
Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Econ. 100 (5), 992–1026.
Boddy, R., Gort, M., 1971. The substitution of capital for capital. Rev. Econ. Stat., 179–188.
Boldrin, M., Woodford, M., 1990. Equilibrium models displaying endogenous fluctuations and chaos: a survey. J. Monet. Econ. 25 (2), 189–222.
Bond, P., Edmans, A., Goldstein, I., 2012. The real effects of financial markets. Annu. Rev. Financ. Econ. 4 (1), 339–360.
Bordalo, P., Gennaioli, N., Shleifer, A., Terry, S.J., 2021. Real credit cycles. Technical report. National Bureau of Economic Research.
Burnside, C., Eichenbaum, M., Rebelo, S., 2016. Understanding booms and busts in housing markets. J. Polit. Econ. 124 (4), 1088–1147.
Caplin, A., Leahy, J., 1994. Business as usual, market crashes, and wisdom after the fact. Am. Econ. Rev., 548–565.
Carceles-Poveda, E., Giannitsarou, C., 2008. Asset pricing with adaptive learning. Rev. Econ. Dyn. 11 (3), 629–651.
Chamley, C.P., 2004. Rational Herds: Economic Models of Social Learning. Cambridge University Press.
Chen, Q., Goldstein, I., Jiang, W., 2007. Price informativeness and investment sensitivity to stock price. Rev. Financ. Stud. 20 (3), 619–650.
Chetty, R., Guren, A., Manoli, D., Weber, A., 2011. Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. Am. Econ. Rev. 101 (3), 471–475.

[27] In the more general case, the optimal policy that maximizes information collection may look more complicated. It remains true, however, that the planner always has an incentive to lean against the market for extreme public beliefs (p_t high or low) as an information cascade occurs and $\overline{F}^{s}_{\theta+\xi}\left(s^*\right)$ goes to 0 or 1.

Christiano, L.J., Ilut, C.L., Motto, R., Rostagno, M., 2008. Monetary Policy and Stock Market Boom-Bust Cycles.
Clarida, R., Gali, J., Gertler, M., 2000. Monetary policy rules and macroeconomic stability: evidence and some theory. Q. J. Econ. 115 (1), 147–180.
David, J.M., Hopenhayn, H.A., Venkateswaran, V., 2016. Information, misallocation, and aggregate productivity. Q. J. Econ. 131 (2), 943–1005.
Eeckhout, J., Lindenlaub, I., 2019. Unemployment cycles. Am. Econ. J. Macroecon. 11 (4), 175–234.
Eusepi, S., Preston, B., 2011. Expectations, learning, and business cycle fluctuations. Am. Econ. Rev. 101 (6), 2844–2872.
Fajgelbaum, P.D., Schaal, E., Taschereau-Dumouchel, M., 2017. Uncertainty traps. Q. J. Econ. 132 (4), 1641–1692.
Galí, J., 2014. Monetary policy and rational asset price bubbles. Am. Econ. Rev. 104 (3), 721–752.
Goldstein, I., Ozdenoren, E., Yuan, K., 2013. Trading frenzies and their impact on real investment. J. Financ. Econ. 109 (2), 566–582.
Golosov, M., Menzio, G., 2020. Agency business cycles. Theor. Econ. 15 (1), 123–158.
Gorton, G., Ordoñez, G., 2019. Good booms, bad booms. J. Eur. Econ. Assoc. 18 (2), 618–665.
Grandmont, J.-M., 1985. On endogenous competitive business cycles. Econometrica, 995–1045.
Guerrón-Quintana, P., Hirano, T., Jinnai, R., 2020. Recurrent bubbles and economic growth. Available at SSRN 3350097.
Guesnerie, R., Woodford, M., 1992. Endogenous fluctuations. Adv. Econ. Theory 2, 289–412.
Harrison, J.M., Kreps, D.M., 1978. Speculative investor behavior in a stock market with heterogeneous expectations. Q. J. Econ. 92 (2), 323–336.
Hirano, T., Yanagawa, N., 2016. Asset bubbles, endogenous growth, and financial frictions. Rev. Econ. Stud. 84 (1), 406–443.
Jaimovich, N., Rebelo, S., 2009. Can news about the future drive the business cycle? Am. Econ. Rev. 99 (4), 1097–1118.
Kaplan, G., Menzio, G., 2016. Shopping externalities and self-fulfilling unemployment fluctuations. J. Polit. Econ. 124 (3), 771–825.
Kindleberger, C.P., 1978. Manias, Panics and Crashes: A History of Financial Crises. Palgrave Macmillan.
Kocherlakota, N.R., 1992. Bubbles and constraints on debt accumulation. J. Econ. Theory 57 (1), 245–256.
Loisel, O., Pommeret, A., Portier, F., 2012. Monetary Policy and Herd Behavior: Leaning Against Bubbles.
Lorenzoni, G., 2009. A theory of demand shocks. Am. Econ. Rev. 99 (5), 2050–2084.
Martin, A., Ventura, J., 2016. Managing credit bubbles. J. Eur. Econ. Assoc. 14 (3), 753–789.
Matsuyama, K., 1999. Growing through cycles. Econometrica 67 (2), 335–347.
Matsuyama, K., 2013. The good, the bad, and the ugly: an inquiry into the causes and nature of credit cycles. Theor. Econ. 8 (3), 623–651.
Miao, J., Wang, P., 2012. Bubbles and total factor productivity. Am. Econ. Rev. 102 (3), 82–87.
Minsky, H.P., 1977. The financial instability hypothesis: an interpretation of Keynes and an alternative to "standard" theory. Challenge 20 (1), 20–27.
Ofek, E., 2002. The valuation and market rationality of Internet stock prices. Oxf. Rev. Econ. Policy 18 (3), 265–287.
Pigou, A., 1927. Industrial Fluctuations.
Scheinkman, J.A., Xiong, W., 2003. Overconfidence and speculative bubbles. J. Polit. Econ. 111 (6), 1183–1220.
Simsek, A., 2013. Belief disagreements and collateral constraints. Econometrica 81 (1), 1–53.
Straub, L., Ulbricht, R., 2017. Endogenous uncertainty and credit crunches. Available at SSRN 2668078.
Straub, L., Ulbricht, R., 2019. Endogenous second moments: a unified approach to fluctuations in risk, dispersion, and uncertainty. J. Econ. Theory 183, 625–660.
Strauss, H., Samkharadze, B., 2011. Ict capital and productivity growth. EIB Pap. 16 (2), 8–28.
Veldkamp, L.L., 2005. Slow boom, sudden crash. J. Econ. Theory 124 (2), 230–257.
Vives, X., 1997. Learning from others: a welfare analysis. Games Econ. Behav. 20 (2), 177–200.
Wen, Y., 1998. Capacity utilization under increasing returns to scale. J. Econ. Theory 81 (1), 7–36.
Xiong, W., 2013. Bubbles, Crises, and Heterogeneous Beliefs. Cambridge University Press, pp. 663–713.